CHAPTER

9

NETWORK OPERATING SYSTEM SERVICES AND APPLICATIONS

After reading this chapter and completing the exercises, you will be able to:

- Identify and understand major network operating system services
- Discuss the different ways that servers run network applications
- Describe the function of monitoring agents
- Specify the functions of the server as a network device

This chapter is about the purpose of having a server. So far, you have been exposed to extensive information about servers, from their environment to identifying and upgrading their hardware, and more. However, none of these things has any meaning except to support the main purpose of having a server (implied in its name): to serve the networking and information-processing needs of an organization. In this chapter, you see how a server uses its various services to make the network run properly and ensure that information gets from one location to the next. An important part of transferring this information is the server's role as a network device, moving data in a fast, efficient manner. Servers also run applications for the network that are too demanding to run on client machines, and they can greatly reduce the processing and storage burden on individual client machines. Because servers are critical to the mission of an organization, administrators must have a way to track server activity and receive immediate notice should something not function properly.

SERVICES

Recall from Chapter 1 that a service provides features to the network. NOS manufacturers include several services with their operating systems, and software packages also often add one or more services. For example, many UPS units include software utilities that install a service to manage the UPS. Many services, including all those discussed in this chapter, include a software interface so that the administrator can configure how the service works. Other services are necessary simply as a function of the operating system. For example, Windows 2000 includes the Plug and Play service as part of the operating system, but you don't administer or configure the service itself. The remainder of this section addresses configurable services.

The services on Linux and UNIX systems are known as **daemons**, but for our purposes we will use only the term "services" unless otherwise necessary.

DHCP

Each host on a TCP/IP network must have a unique IP address. If two hosts on the network have the same IP address, communication problems occur. In Windows NT, the first host that enters the network with a given IP address retains it when a second Windows NT host with the same IP address enters the network. The second host, however, cannot communicate. Both hosts enter a record of the duplicate IP address problem in the System log. Duplicate IP addresses are one of the most common problems in IP networks with **static IP addresses**—that is, where someone manually enters a permanent IP address for a host, which includes IP configuration such as the subnet mask, name servers, and routers. Human error and a lack of good records can lead either to wrongly configured IP addresses that do not communicate properly on the network or to duplicate records (as just mentioned). Furthermore, it requires a significant investment in time for technicians to manually enter static IP addresses on hundreds or thousands of individual clients.

This kind of problem can be avoided by the use of a **Dynamic Host Configuration Protocol (DHCP)** server. A **DHCP server** automatically allocates IP addresses to hosts on the network. Recipients of DHCP-allocated IP addresses are known as DHCP clients, and in order to receive the IP address, you must specify this in the client's network properties. For example, Figure 9-1 shows the NETCONFIG utility on a Red Hat Linux server, which you use to enter IP configuration information. The setting "Use dynamic IP configuration (BOOTP/DHCP)" instructs the client to seek a DHCP server on the network from which to obtain an address. (**BOOTP** is the Bootstrap Protocol and, similar to DHCP, uses a BOOTP server that can distribute IP addresses to clients.)

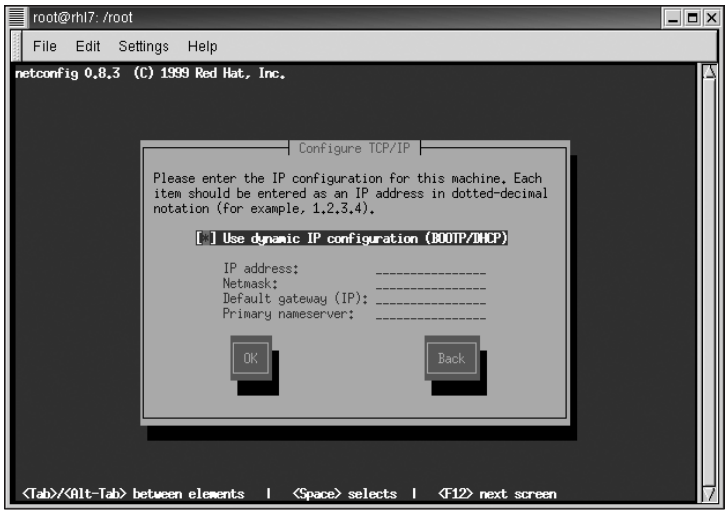

Figure 9-1 "Use dynamic IP configuration (BOOTP/DHCP)" means to use a DHCP server

9

The Lease Process

When a client receives an IP address from the server, it has obtained a **lease**. Similar to the lease of property or a car, this lease is for a limited duration. Administrators can adjust the lease from a few hours to several days, depending on the nature of client access. For example, in a typical office building where the network hosts are mostly desktop computers and very few changes occur on the network, you might configure leases that last for days. However, if you run an ISP (Internet service provider), the lease duration is probably an hour or less, because you do not want users who have disconnected to retain the address lease while it would otherwise be available to the next user.

The client initiates the lease process with a **discovery broadcast** seeking a DHCP server. A broadcast to all hosts is necessary, because the host does not yet have an IP address with which to communicate, nor does it know the IP address of the DHCP server. All available DHCP servers respond to the client with an **offer** of a specific IP address and configuration. The client responds with a **request**, which is an acceptance of the offer from one of the randomly selected DHCP servers while rejecting offers from the rest. Finally, the server issues the IP configuration, which it confirms with an **acknowledgment**. Using the first letter of each of the four steps, you can easily remember the process with the mnemonic "DORA." Figure 9-2 illustrates this process.

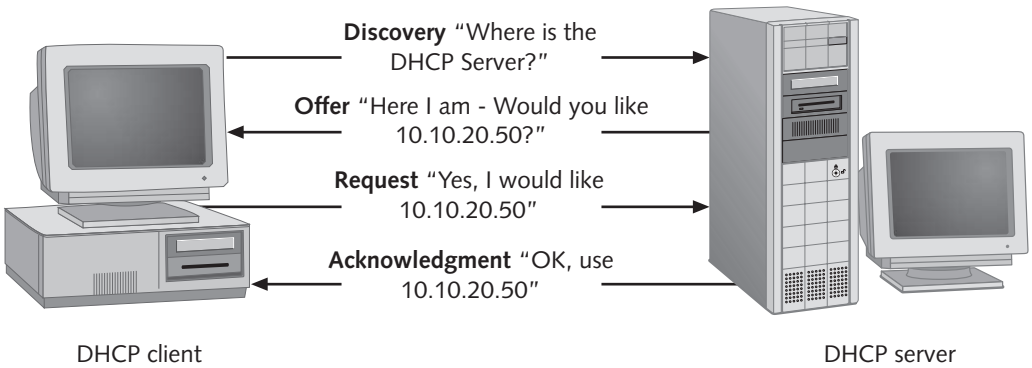

Figure 9-2 The DORA lease process

Although using DHCP for client IP address configuration helps ensure accurate IP addressing and reduces the administrative burden of manual configuration, it is best not to use DHCP for servers, printers, and network equipment so that those addresses are always consistent. You do not want a printer, for example, to lease a changing IP address because it becomes unavailable to clients attempting to access the printer's original address.

A DHCP server usually distributes not only the IP address and subnet mask to the client, but also many other configuration items such as the IP address of the gateway (router) and DNS (Domain Name System) and WINS (Windows Internet Naming Service) name servers. In Windows 2000 or NT, you can configure DHCP by opening the DHCP item (Start, Programs, Administrative Tools, DHCP). NetWare also has a GUI console that you can use to configure both DHCP and DNS. Figure 9-3 shows the NetWare DNS/DHCP Management Console interface. Linux, because of its UNIX origins, still requires configuration of the DHCP service through the text-based dhcpd.conf file.

When configuring DHCP, you must know the range of IP addresses you want to use on the network (known as a **scope**). You can use private IP addresses, which are practically limitless and for which there is no direct access from Internet hosts, or you can use public IP addresses. As a security precaution, most administrators prefer private IP addresses for their network hosts. Once you determine the scope you want to use, you enter a starting IP address and an ending IP address on the DHCP server. Hosts then lease the IP addresses one at a time in sequential order. The administrator ensures that enough IP addresses are available in the scope to service all hosts on the network, and usually add a margin for growth.

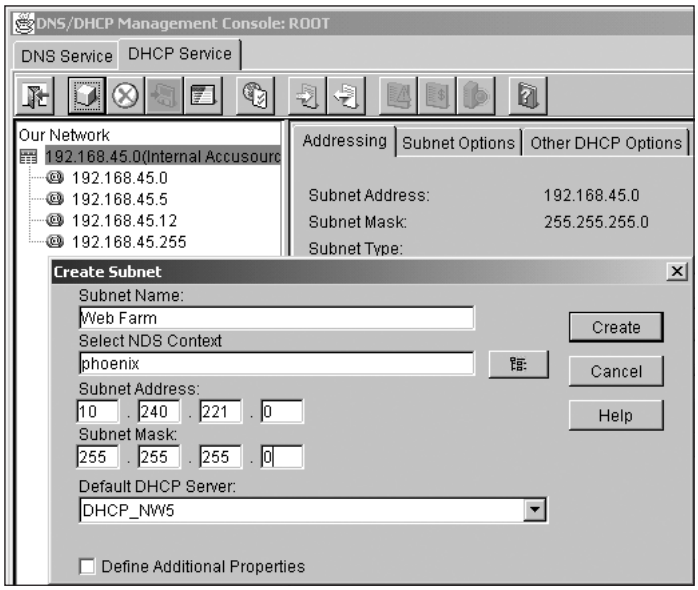

Figure 9-3 Use the NetWare 5.1 DNS/DHCP Management Console to configure DHCP

Be careful about testing IP scopes on the network. If someone introduces a DHCP server with incompatible or duplicate addresses on the network and the organization is large, then the new DHCP server will start to answer DHCP client requests. Those clients will then be incorrectly configured in the network, and it could be very difficult to find out exactly where the rogue DHCP server is and shut it down. Windows 2000 requires administrators to perform an additional step of authorizing a DHCP server before it functions on the network. This helps to ensure that a rogue DHCP server does not enter the network.

For sake of network availability and redundancy, implement more than one DHCP server. If one server fails, another DHCP server can continue to provide service. The redundant DHCP server(s) will have a compatible range of IP addresses usable on the network, but none of the IP addresses overlap from one DHCP server to the next. If it is not practical to use two DHCP servers in the same network segment, you can program routers to forward BOOTP or DHCP broadcasts to another DHCP server on a different segment. Also, the operating system might include a mechanism by which a host listens to DHCP discovery broadcasts and forwards the discovery inside a packet that is sent directly to the IP address of a DHCP server. For example, in Windows NT and Windows 2000, the DHCP relay agent performs this function.

Some organizations consider use of a DHCP server a security threat. The concern is that a stranger could enter the building as a visitor and locate an available network jack. Then, by turning on a laptop and leasing an address, the stranger's laptop becomes a part of the network and can begin to explore the network to provide information that would allow unauthorized access to resources.

DNS

People tend to remember names better than numbers. However, network hosts are identified using numbers. Though administrators are likely to have memorized the IP addresses of several key network resources such as printers, routers, web servers, mail servers, and logon servers, nobody memorizes the IP addresses of an entire enterprise. When users want to access a network resource or web site, they do not type in an IP address such as *www.199.227.124.246.net*. Instead, users would type the name, *www.accusource.net*.

It is the DNS server that transparently allows users to operate this way, because the **Domain Name System (DNS)** server stores a record of both the IP address and host name, and uses these records to service name resolution requests. DNS is actually a replacement for the HOSTS plain text file used in all major NOSs. The **HOSTS file** contains static, manual entries of host-to-IP address mappings, much like this entry for a computer named websrv6:

109.54.94.197 websrv6.accusource.net

With the growth of the Internet, using a HOSTS file to resolve web sites became impractical years ago. DNS also requires manual configuration but is much more flexible in its administration. However, many organizations still use the HOSTS file because at boot time, the file is loaded into memory for quick resolution of a network host without a DNS server.

DNS Name Resolution

Although the length of time you wait for most host name resolution requests to resolve is perceptibly short, the resolution process might actually take place through several DNS servers across a wide geographic area. For example, when you type *www.accusource.net* into a web browser, the computer issues a query to its configured DNS server, which then resolves the name to an IP address. As the Internet is large and even the most capable DNS server does not have all DNS records, the DNS server often references other DNS servers to answer a request. Forwarding the request like this is known as a **recursive query**. While waiting for another DNS server to resolve the request, the original DNS server becomes the resolver while the client simply waits for an answer. In a recursive query, the resolution burden rests with the DNS server to either resolve the name itself or resolve it using other DNS servers (compare this to an iterative query, following). This process is illustrated in Figure 9-4.

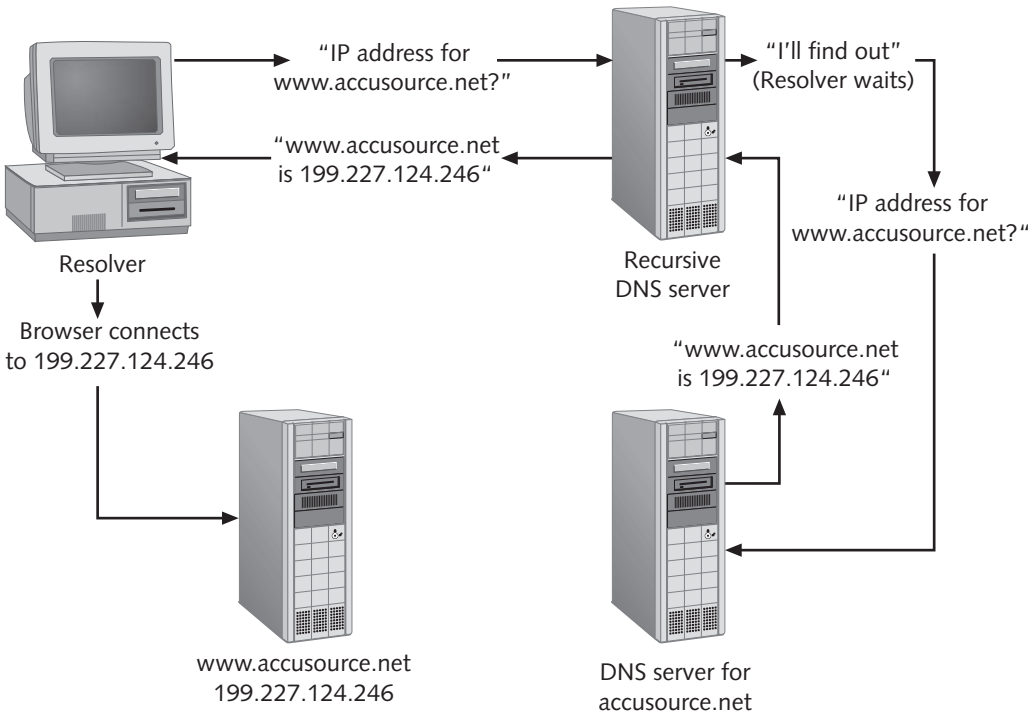

Figure 9-4 A recursive query

Internic is the governing body responsible for naming and managing the DNS domain namespace. The domain space is a hierarchical tree of names created to manage DNS requests using separate authoritative DNS servers to manage each domain space such as .com, .gov, .edu, .net, .org, and so on. See more about Internic at *www.internic.org*.

The DNS server can also be configured to work in an **iterative query** mode. When the DNS server receives a request for which it does not have an answer, it refers the resolver to more authoritative DNS servers further up the hierarchical DNS namespace tree. The iterative DNS server is then relieved of the resolution burden, because it is the client resolver that references other DNS servers. In turn, the next DNS server to which the resolver passes the request might forward the request to yet another DNS server.

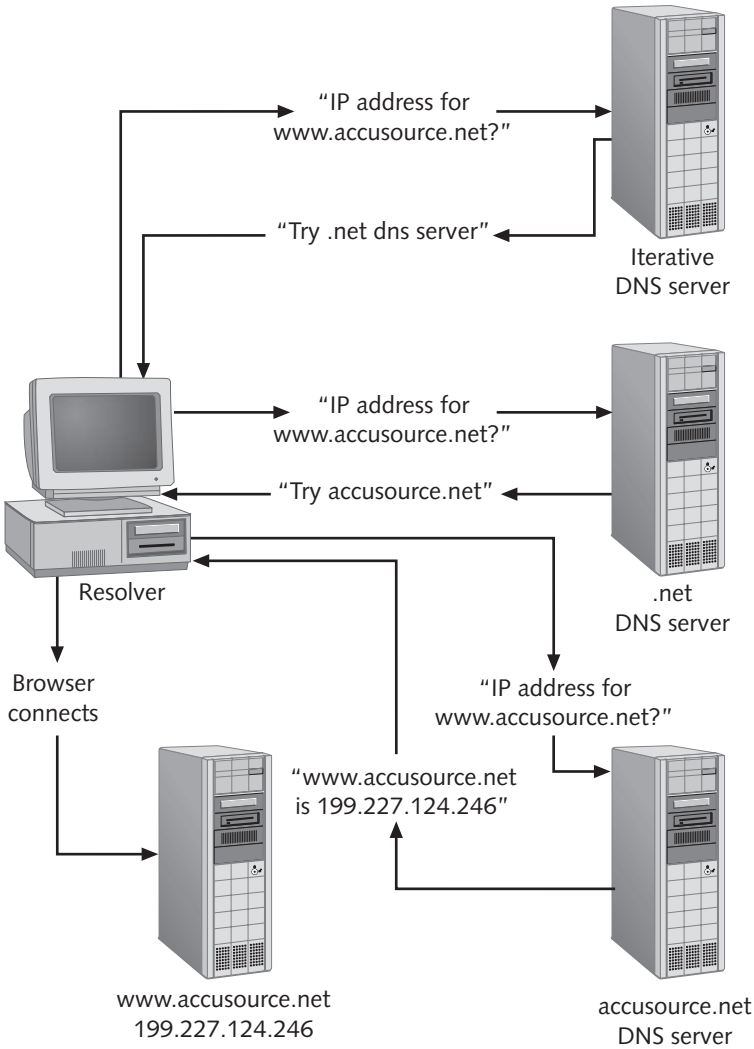

Figure 9-5 An iterative query

For infrequently accessed Internet sites or sites that have only recently become registered, a client web browser might time-out on the first attempt while waiting for various DNS servers to resolve the request. Often, if you try a couple more times, DNS servers through which the request passed the first time will have cached some results, helping to speed up the resolution.

DNS Zones

Many large organizations have one or more full-time employees whose main function is to create and maintain DNS database records. Except for Dynamic DNS records, which this chapter discusses later, most DNS records require someone to make a manual entry on the DNS server. Each NOS has a different interface for creating such records, but if you know DNS, then you can create the DNS records on almost any GUI once you get used to its interface. For example, the DNS/DHCP Management Console for NetWare 5.1 has a handy toolbar for configuring DNS (see Figure 9-6).

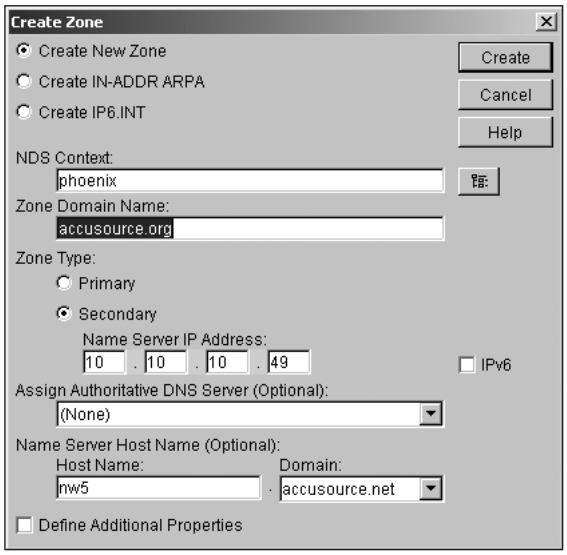

Figure 9-6 Use the DNS/DHCP Management Console to create DNS servers and zones

This section provides some basics and helps you learn some of the most commonly used types of records; the rest is a matter of learning the interface.

Before you go about creating DNS records, you need a DNS zone. The **DNS zone** is a naming boundary that you probably see every day when using the Internet. For example, in *www.course.com*, "course.com" is a zone and "www" is the name of the web server in that zone. The zone "contains" the records. When you create a zone, you also add a record for the DNS server that is the **Start of Authority (SOA) server**. The SOA server is the authoritative source for information about the domain, and the domain cannot function without it. In the configuration screen for a new zone, the NOS might include fields for you to provide the name SOA record as in Figure 9-6 (referenced above) where the NetWare console labels it "Assign Authoritative DNS Server." Windows 2000 automatically assumes that the SOA server is the DNS server on which you create the zone, and creates it for you. If this is not correct, you can manually change this. An SOA record is also required for the reverse lookup zone (discussed later in this chapter).

If you use UNIX/Linux, then you need to learn the text-based format in which you add zones to the named.conf file, and then create separate database files (*.db) that store the individual records for each zone. One of the major flavors of UNIX is BIND (Berkeley Internet Name Domain), and it has a very reliable, time-tested facility for DNS. In fact, Internet DNS servers are typically UNIX servers. You can use BIND DNS on a Linux server; however, although the service is very reliable, the interface is text-based and not as intuitive as other NOSs. For example, here is an excerpt from a named.conf file for accusource.net:

```
zone "accusource.net" {
type master;
file "accusource.db";
```

After you create the zone, you should also create the **reverse lookup zone**, otherwise known as the IN-ADDR.ARPA zone. This zone performs the reverse of a normal query: Instead of resolving a name to an IP address, you're resolving an IP address to a name. A reverse lookup is not as frequently used as a forward lookup, but it is useful if you know the IP address of a host and want to determine its name. Recently, I suspected that a server with a particular IP address was returning errors, but I couldn't remember which server had that specific IP address. By using the *PING* -a command, which returns the host name of the associated IP address, I saw the host name of the server. The host name implied its physical location and role, making it easier to find the specific server. Reverse lookups are also useful for verifying the claimed identity of a connecting host as a security precaution. Another reverse lookup tool is the NSLOOKUP tool available with Windows NT/2000.

DNS Records

Creating records is a straightforward process that again is mostly a matter of using the operating system interface or manually entering records in UNIX/Linux.

The following list is not exhaustive, but it itemizes the main DNS records of which you should be aware:

- **SOA record**. As discussed earlier, an SOA record identifies the authoritative name server for a given domain.
- **Name Server (NS) record**. Specifies what DNS servers are delegated servers for the domain, meaning that the server specified in the record can resolve queries authoritatively.
- **PTR (pointer) record**. The actual record used in reverse lookups.
- **Address (A) record**. Also known as a host record, this is the actual record that resolves the host name to the IP address.
- **Mail Exchanger (MX) record**. Routes mail to the appropriate server(s) for members of the domain. If you have multiple MX records for the same domain, you can prioritize the mail servers with numbers between 0 and

65,535, with lower numbers having the highest priority. In Figure 9-7, this mail server with a priority of 10 would have a higher priority than a mail server with a priority of 20.

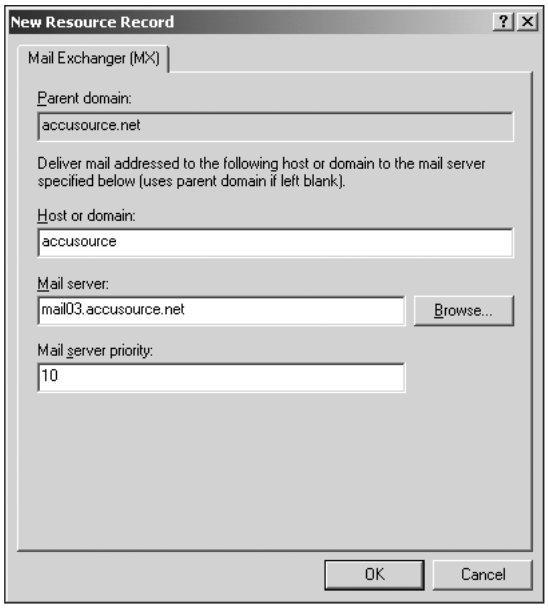

Figure 9-7 A Windows 2000 MX record with a priority of 10

- **CNAME record**. Stands for canonical name record, and is an alias that points to another host. You use CNAME records every time you browse the Internet. When you type *www.xyz.com* into a web browser, "www" represents a server located in the xyz.com domain. However, very few companies actually name their web servers "www." Let's say that the xyz.com domain's web server is actually named websrv13. That would make its true Internet address websrv13.xyz.com. However, most organizations prefer not to expose the true name of their web server on the Internet for security reasons. The xyz.com DNS administrator would then create a CNAME record that directs inbound requests for www.xyz.com to websrv13.xyz.com in a completely transparent way to the users. CNAME records are also useful in directing inbound requests to any of multiple servers. For example, if xyz.com has 13 web servers, a CNAME for each server can be used to direct inbound requests to any one of the 13 web servers.

A common way of directing inbound requests to any of multiple web servers is to use the round-robin method, which you configure using the DNS server properties. DNS directs the inbound request to a host, and then sends the host to the bottom of the list. DNS directs the next request to the next host, and in turn sends that host to the bottom of the list and so on, until the original host returns to the top of the list.

Be meticulous about DNS administration. DNS is a critical service because it points to other services. For example, it might direct network clients to mail servers, logon servers, web/intranet servers, and so on. DNS problems can bring down an entire network. In fact, some major web sites have become unavailable for hours or days because of an incorrect DNS record.

Before closing the discussion on DNS records, I should point out an exciting new technology known as **Dynamic DNS (DDNS)** that greatly reduces the burden of manually creating DNS records. In network environments where Microsoft operating systems exist, NetBIOS names, instead of host names, identify network nodes. Introducing the mechanisms necessary to use and resolve NetBIOS names adds overhead and administration to the network. Beginning with NetWare 5.1 and Windows 2000, you can now rid the network of reliance on NetBIOS because you can instead use DDNS. Using traditional DNS, this would have required the administrative nightmare of manually creating, updating, and deleting DNS records, which usually makes registering hundreds or thousands of desktop clients completely impractical. Now, because of an interaction with the DHCP service, DDNS can accept automatic registrations from clients when they receive their IP configuration from the DHCP server. This capability is not limited to NetWare 5.1 and Windows 2000—any environment including UNIX that implements RFC 2136 can use DDNS. (RFC is "request for comments" and is the titling method used by the Internet Engineering Task Force to identify documents.)

Types of DNS Servers

There are three primary types of DNS servers: primary domain servers, secondary domain servers, and caching-only servers.

- **Primary domain server** is the starting point of all DNS records, containing a read/write-capable zone database: you can add, remove, or modify DNS records from a primary domain server.
- **Secondary domain servers** receive a read-only copy of the zone database from a primary domain server. Secondary domain servers are useful for providing redundancy and load balancing.
- **Caching-only servers** have no zone database, either of their own or copied through a zone transfer from a primary domain server. Caching-only servers mostly function only to improve performance by reducing the number of forwarded queries. When a caching-only server retrieves a resolved query, it stores the result in memory for a period of time known as the **time to live (TTL)**.

When a DNS zone is created, there is a master-slave relationship between the primary master server, master servers, and slave servers as follows:

- **Primary master servers** are the first and final authority for all hosts in their domain. There is only one primary master server per zone, and they are the source for records that are copied to master or slave DNS servers.

- **Master servers** are authoritative DNS servers that transfer zone data to one or more slave servers. ("Authoritative" means that the server is configured to host the zone and return query results.)
- **Slave servers** are authoritative servers that receive the **zone transfer** (a copy of the zone DNS database) from the master server and are named in the zone by NS records.

WINS

The **Windows Internet Naming Service (WINS)**, as the name implies, is a Microsoft invention. In a Microsoft network, there are two types of names, and a client might possess one or both of them: a host name, as discussed in the previous section on DNS, and a NetBIOS name. WINS is a service that performs a similar function to DNS, except that it resolves NetBIOS names instead of host names. (Recall that **NetBIOS** is a broadcast-based name resolution scheme where a client simply broadcasts the NetBIOS name of the computer it wishes to reach to all of the computers on a subnet.) The broadcast message identifies a computer that acknowledges the broadcast and establishes a communication link. The limitations of this are clear from the fact that it is a broadcast message and therefore cannot be routed to another subnet. This resolution method also creates more traffic on the subnet due to the nature of broadcast messages. WINS mitigates both of these problems by registering the NetBIOS names and resolving them to IP addresses that are routable and destination-specific. WINS is only a necessary evil to overcome the inherent limitations of NetBIOS names. Some network applications use WINS to locate network hosts. Also, mapping drives using the *NET USE* command use WINS.

Windows 2000 changes this reliance on WINS for name resolution by registering Windows 2000 computers, servers, and domain controllers via DDNS (as discussed earlier in this chapter). In a strictly Windows 2000 networking environment, DNS is the only name resolution service that is necessary. However, if you have pre-Windows 2000 Microsoft clients, or if you have network applications that require a NetBIOS interface to resolve host names, you will likely need to run the Windows NT or Windows 2000 WINS service.

WINS is a method by which user-friendly NetBIOS computer names can be resolved to IP addresses for the purpose of allowing such communication between hosts on different subnets. In order for this resolution to occur, network hosts must have a way to dynamically add, remove, or update their names in the WINS database. The WINS database is only as good as the accuracy of the records that are contained in it. The database must reflect any changes made to a client's configuration with as little delay as possible. The administrative overhead required to make these changes manually would be prohibitive, so the need for a dynamically updated database is clear.

Name registration occurs when a WINS client requests the use of a NetBIOS name from the WINS server. The WINS server can either accept or reject the request for a NetBIOS name made by the WINS client. The response that is given depends on several factors. If the requested name does not already exist in the WINS database, the WINS server

accepts the request and creates a record containing the name of the NetBIOS client and its IP address, among other things. The WINS server also sends an acceptance message containing the TTL parameter to the requesting client.

Resolving NetBIOS Names

Windows 2000 client name resolution follows a very specific order. WINS is one step in that order and is used only after exhausting the first two options. When a Windows 2000 client attempts to reach another host by name, the client determines whether or not to use DNS. DNS resolves the name for the client if there are more than 15 characters or if there are periods in the name. NetBIOS names must be 15 characters or less and cannot contain periods, making DNS the only service that could properly resolve the name.

If the name is 15 characters or less and does not contain periods, the second step is for the client to check its remote name cache stored in RAM for resolution to the IP address. (This is the first step for non-Windows 2000 clients.) The remote name cache is a location in memory that stores recent NetBIOS names to IP address resolutions. This is the most efficient method of name resolution because all the client reads is its own memory.

If the name is not in the cache, the client will contact the configured (either statically or through DHCP) WINS servers to see if there is a record in the database for the requested computer name. If there is a record for the NetBIOS name, the WINS server returns the IP address to the requesting client. The client then uses the address to connect to the desired server.

If the name is not in the WINS database, the client will broadcast the name to the subnet in hope that the desired server will respond. The main limitation of broadcasts is that routers cannot forward broadcast messages to another subnet—so a client that is on another subnet will not be able to respond.

If the broadcast fails, the client will check its LMHOSTS file to see if there is a static entry for the NetBIOS name. The **LMHOSTS file** is a static text file that lists NetBIOS-name-to-IP-address mapping and is the NetBIOS equivalent of the HOSTS file. You can configure entries in the LMHOSTS file to load individual records into the remote name cache at system startup. Therefore, Windows 2000 queries those LMHOSTS records at the very beginning of this process.

If the LMHOSTS file fails to provide resolution, the client queries the DNS resources. First, it queries for HOSTS files. For example, if you know that a client will need to access a specific logon server, adding the name of the host separated by the host's IP address automatically loads the name mapping into memory. If this fails, the client queries the configured (static or DHCP) DNS servers. Figure 9-8 illustrates the process of NetBIOS name resolution.

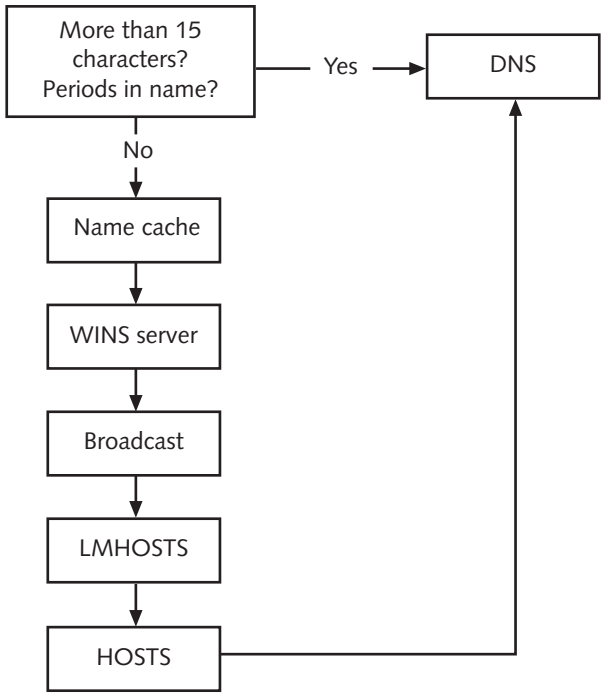

Figure 9-8 NetBIOS name resolution order

As a memory aid, you can take the first letter of each step in the order and correlate it to the first letter of the following sentence: "Can We Buy Large Hard Drives" (Can = Cache, We = WINS, Buy = Broadcast, and so on).

WINS Replication

As a network becomes larger than just one subnet, and for the sake of redundancy, you should use more than one WINS server. No matter how many WINS servers you have in your network, you will still only have a single WINS database. With multiple servers and one database, **replication** (copying the database from one server to another) is extremely important to ensure that all of the WINS servers have a consistent copy of the database.

You configure each WINS server with one or more replication partners in such a way that all of the WINS servers will eventually receive each change that occurs on a single WINS server. When a WINS client registers its name and IP address with a WINS server, that server is the owner of the record and the record propagates to all of the other WINS servers through a series of replication partners. WINS replication is incremental, meaning that only the changes in the database replicate, not the entire database. The entire database replicates only when you install a new WINS server on the network.

All WINS servers maintain an owner-version mapping table that it builds dynamically and stores in memory. WINS servers use this table in the replication process to identify the server's replication partners and the version of the information that is contained on that server. The field in the table that is important for replication is the Highest ID field, which stores the highest known version ID that is contained on the replication partner. Remember that all records have a Version ID field, which records the highest value contained in all of the version IDs received from a particular replication partner.

Pull Replication Partners

A **pull replication partner** is a WINS replication partner that requests and then accepts changes from its push replication partners. At specified intervals, and when a WINS server starts, the pull replication partner requests all records from the push partner that have a higher version number than the last entry received from that particular pull partner. The amount of time between pull requests is called the **replication interval**, which can be set globally for all WINS servers or can be set at each WINS server that is configured as a pull partner.

The pull replication partner initiates the replication by sending a **replication trigger** message to a push replication partner. When the push replication partner receives the trigger, it scans for the highest version ID contained in the database for all WINS servers on the network. This highest ID for each owner is compared to the highest ID records that the pull replication partner currently has. All of the records that have version ID numbers that fall between the two replicate to the pull partner initiating the replication.

Push Replication Partners

A **push replication partner** is a WINS replication partner that responds to requests for changes from its pull replication partners. Push replication occurs when the WINS server starts or when a set number of name-to-address changes have occurred in the replica contained in the push partner. You can set this value in the properties of the push replication partner WINS server. By default, the number of changes that are required to initiate a push replication is 20. If you enable persistent connections, which is not the default, the number of changes that are required to initiate a push replication is set to zero and any changes are replicated as soon as they are made. You can configure a push replication partner in this situation to initiate changes less frequently. The push/pull replication process is illustrated in Figure 9-9.

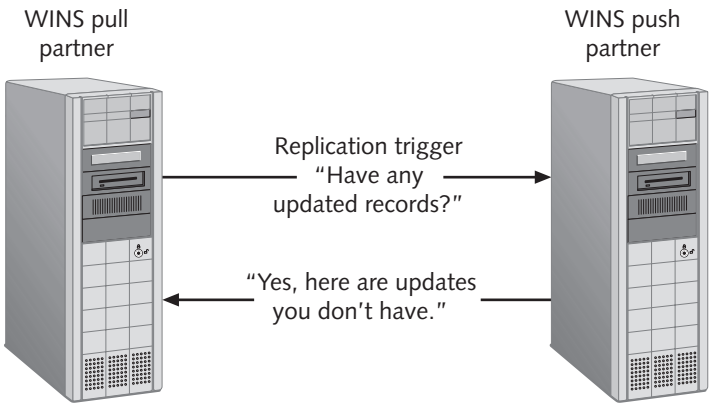

Figure 9-9 The WINS replication process

Push/Pull Replication Partner

A push/pull replication partner acts as both a push and a pull replication partner. This is the default replication functionality for Windows 2000 WINS servers. In general, setting up WINS servers to be push/pull replication partners is the simplest and most effective way to ensure full replication between replication partners.

Avoid the use of unidirectional replication partners. There are some cases in large networks where the use of unidirectional partners can limit the amount of traffic created over slow WAN links. When configuring unidirectional partners, you should be careful that each server has at least one replication partner. Also, balance unidirectional partners configured over a WAN link with another link in the opposite direction to some other location in the network. Finally, configure primary and secondary WINS servers as direct replication partners of each other.

WINS Proxy Agent

Non-WINS clients use broadcast messages to locate other nodes on the network by default. Broadcast messages are not routed, so that if the client attempts to contact another computer located on another subnet, the non-WINS client must use a statically configured LMHOSTS file to resolve the IP address so that the client can be contacted.

A **WINS proxy agent** is a WINS-enabled computer that you configure to listen on the subnet for WINS broadcast messages, such as query, refresh, release, and registration. The WINS proxy then communicates with the WINS server to resolve or register NetBIOS names.

When a non-WINS client sends a broadcast query for a node on the network using a NetBIOS name, the node with that name responds if it is on the same subnet as the broadcasting node. The WINS proxy agent does not respond to a query message if the

non-WINS client is on the same subnet. The WINS proxy agent determines that it is on the same subnet as the requesting client by comparing its address with the address of the requesting client using the subnet mask.

If the node is not on the same subnet, the WINS proxy agent intercepts the broadcast and checks its own cache to see if it has a record of the NetBIOS name and its associated IP address. If it is in the cache, the WINS proxy agent sends the IP address to the non-WINS client so that it may now send routable, directed packets to the desired node.

If the name is not in the cache, the WINS proxy agent queries the WINS server using directed packets. The WINS server then responds with the IP address that is associated with the NetBIOS name. The WINS proxy agent then forwards the information to the non-WINS client and stores the information in its cache for future use.

To configure a WINS-enabled Windows 2000 computer, you must edit its Registry by adding the value EnableProxy to the HKEY_LOCAL_MACHINE Registry subkey as follows:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netbt\Parameters.

You must set the value to 1 to enable the WINS proxy agent. The type of data is REG_DWORD, and is shown in Figure 9-10.

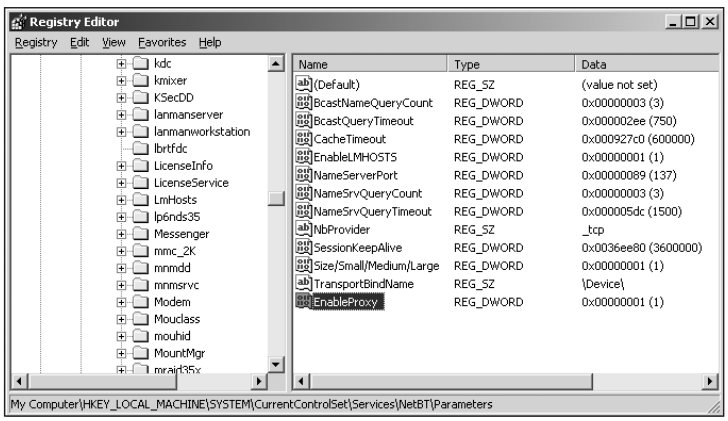

Figure 9-10 The EnableProxy registry subkey

Mail Services

Email used to be available only for internal communications within the office. However, with the ever-expanding Internet, email has become a global medium for both internal and external communications. Email is now possibly the most heavily utilized service in most organizations, and is also the most visible to users. If the web site goes down, users might not notice for a while and can probably continue most of their job functions. If

email services fail, users start calling immediately. Running email services involves knowing the main components, adequately preparing the hardware, and planning for disaster, which is discussed in Chapter 10.

Email Protocols

Email requires a transport protocol to get it from one place to another. Each email protocol has its own unique purpose, as listed below:

- **Simple Mail Transport Protocol (SMTP)**. As the name implies, this protocol is very simple. In fact, by itself it only transports basic text—you could not use SMTP alone to send a binary email attachment such as a multimedia file. MIME (see next) adds this functionality. SMTP is the protocol that transfers or forwards mail to an email server. However, when clients retrieve email, they typically use IMAP or POP3.
- **Multipurpose Internet Mail Extensions (MIME)**. This protocol adds the mail capability of attaching and transferring multimedia file attachments. To use MIME, you must also have an email client capable of decoding the MIME format.
- **Post Office Protocol 3 (POP3)**. Both POP and IMAP (see next) use the SMTP protocol for the actual transfer of information, and both allow messages to be stored on the mail server for incoming email. Then when you log on, you download all the email using POP3 with no selectivity.
- **Internet Message Access Protocol (IMAP)**. In its current version (IMAP4), IMAP allows the email client to leave messages on the mail server even after logging on instead of downloading each one. This is useful for keeping email in a central location where it can be organized, archived, and made available to remote locations. You can also search through mail for certain keywords while it is still on the mail server, and selectively download messages that match the search. IMAP integrates with MIME so that users can read the mail header information and then decide whether to separately download the attached files. Compare this to having to download the files with the email before you can read the header, which can take considerably longer with large files.

Notice that some of the email protocols depend upon or interact with others. The protocols are not mutually exclusive, and one email transaction might require several protocols. For example, a friend might send you an email with an attached JPEG picture from his vacation. The message is sent to the server using SMTP, and MIME allows the JPEG picture as an attachment. You dial up to an ISP with IMAP capability and see the header that reads: "A great picture from my vacation." Because you are on a very slow dial-up line and expect the JPEG to be large, you decide not to download it. However, you are expecting an important JPEG file from a co-worker that you want to use in a presentation the next day—so you search for it, again because of IMAP capability. IMAP allows you to find the message, which can then be downloaded.

All email applications include a message store that stores mail until the client downloads it.

Email Server Applications and Requirements

Email has in many ways reduced paper communications, such as the ubiquitous interoffice memo. Software products such as Novell GroupWise and Microsoft Exchange extend the basic communication features of email, offering a host of features such as online collaboration, calendars, newsgroups, contact information, chat functions, and more. Mail servers must also be able to store the messages until the client requests and downloads them. These functions significantly contribute to the hardware requirements of a mail server, which vary broadly depending on the product and utilization. For example, Linux can use the BIND Sendmail program, which requires light hardware resources, because it offers primarily plain email functionality and does not include many of the shiny new features found in some other email server products. Nevertheless, Sendmail is a long-standing favorite among email administrators as it is highly reliable. Find out more about Sendmail at *www.sendmail.org*.

Because a mail server is highly visible, it must be protected with regular backups and additional servers for purposes of redundancy and failover. Also make sure that additional DNS servers are available. Outgoing mail from your organization could go to any of countless other domains, and if you are suddenly without a DNS server to instruct the mail server as to the location of these other mail servers, the mail will not go.

Consider creating an alliance with another organization so that you can act as fault-tolerant partners. If one organization's site goes down and is unavailable to receive mail, the other organization will receive and temporarily store mail until the other's site is back up. Otherwise, the email is likely to bounce back to the sender.

Prepare other servers and network functions in the organization to send and receive mail. The following list should be a minimum starting point:

- Most organizations have firewalls to protect the network from the Internet. Be sure to open the firewall to the protocols mentioned in the previous section so that messages can pass into and out of your organization.
- Add the appropriate MX records to identify and prioritize the email servers. If you also have a backup remote site with mail servers, be sure to configure the MX priority number to be higher than the local mail servers.
- Provide plenty of hardware. If there's any server that you can't underestimate, it's the mail server. Mail taxes every main hardware component, including memory, processors, and hard disk storage. For example, Microsoft Exchange 2000 Server requires 128 MB of RAM (256 MB recommended), 500 MB on the hard disk where you install Exchange 2000, and 200 MB of additional

space on the system drive. Note that this is on top of the Windows 2000 operating system, which has minimum requirements of its own, and that a server configured with these minimums will only be able to handle very light mail traffic. A production mail server is likely to have upward of 1 GB of RAM and several gigabytes of disk space to store all the messages.

Mail programs allow the administrator to set quotas on the user mailbox size. I strongly recommend mail quotas to conserve disk space and improve backup and restore time. A large credit card company in my area made the mistake of not implementing quotas, and many users have mailboxes over 1 GB in size. The strain is so heavy on the mail system that about every week a hard disk fails. Figure 9-11 shows the interface for configuring mail quotas for Exchange 2000.

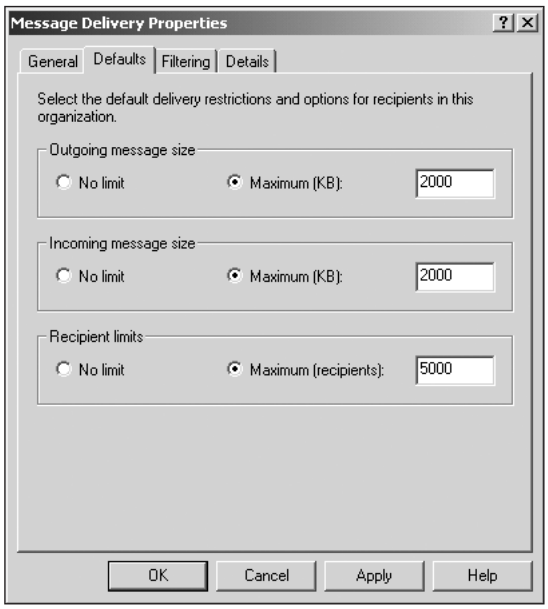

Figure 9-11 Configure limits (quotas) on the size of messages and the permitted number of recipients

- Mail server hard disk performance quickly degrades because of the constant write activity. Be sure to defragment the hard disk regularly.
- Mail is also very taxing to network bandwidth. Make sure that you have plenty of available bandwidth, and also consider using a multihomed NIC configuration (such as adapter teaming).
- Because mail service is so active and demanding, mail servers should generally be dedicated to mail instead of serving multiple purposes. If necessary, you might include a relatively light additional service such as DHCP, but also

using the mail server as a heavily utilized database application server would be out of the question.

- Email is often the primary vehicle by which viruses invade a network. It would be foolish not to implement a comprehensive and reliable antivirus solution to check email.

Choose your mail server product very carefully. If you later decide to change to a different product, it can be very disruptive, even though most products include a means by which a competitor's mail service can be converted. Also, the mail product can alter the operating system. For example, Exchange 2000 adds modifications to the Windows 2000 Active Directory schema that are permanent. Even if you uninstall Exchange 2000, the schema changes remain. On the plus side, the mail product can greatly simplify administration. For example, you might be able to create a user account and simultaneously create a mail account instead of creating them separately. Likewise, carefully select the email client to reduce licensing expenses and support burdens.

- Prepare for legal issues. Most companies consider employee email to be company property when delivered to a company address using company equipment. Many employees disagree. When users are hired, be sure to notify them in writing that email is neither the property of the employee nor private, if that is the company position.

Web and FTP Servers

Even small and medium-sized organizations usually have a web presence—a web site seems to be as necessary as company stationery. A few years ago, some segments of the IT talent pool knew how to set up a web site, usually on a UNIX server. Now, it is a given that administrators know at least the basics of web site configuration and administration. This section addresses the server's role as a web and FTP server.

Web Server Requirements

As with mail servers, web and FTP servers (which we will collectively refer to as only "web servers") are normally dedicated solely to those respective functions. The specific type of web server you run depends largely on the needs of the organization, the number of hits (visitors) on the site, and the type of access. For light-duty web service, virtually any server might be able to run the site. An administrator might configure Internet Information Server (IIS) on a Windows NT server, create some basic content, and probably have the site running and available in less than an hour. However, a web site of this scale will not suit most active business concerns. A medium-sized or large organization usually requires web service with:

- Dedicated bandwidth to Internet traffic that is separate from the LAN connection to the Internet. Although you could share the bandwidth, LAN traffic might adversely affect the responsiveness of the web site and vice versa.

- A large, fast, and redundant disk storage solution such as Fibre Channel RAID for multimedia or file downloads. Consider placing download files on an FTP server because it is a more efficient file transfer protocol than HTTP. You can configure the link on the web page to point to an FTP site to initiate the download—the user does not have to type in an FTP address. Plain web page content by itself does not usually require a great deal of hard disk space.
- Adequate processing power. Processing power requirements vary greatly. For example, a web server can answer thousands of hits per day, yet have relatively little processor utilization if the web content is simple. However, more complex functions such as online transaction processing (OLTP), as required for Internet credit card purchases, can require significantly more processing power. The transaction itself will take place in a back-end database on a different server; however, the encryption required to validate the server and clients requires processing power. That's why if you visit your bank's web site, the home page probably loads quickly (little processing involved). But logging in using 128-bit Secure Sockets Layer (SSL) encryption to access your bank account requires a lot more processing power, and it may take several seconds to access your information.
- A digital certificate that validates the organization to the public. The digital certificate is an authentication mechanism using a form of digital identification, which enables SSL encryption between clients and hosts. For example, VeriSign (*www.verisign.com*) is a leading source of digital certificates. When visitors access a site, they want to know that the site is reputable and trustworthy, especially for online SSL transactions. A certificate is like a recommendation from the Better Business Bureau, only better.
- Adequate memory. As much of the entire web page content as possible (except for download files and streaming media) should fit into main memory. As visitors access the site, the pages load into cache, where they are available at a much faster rate than by using disk retrieval.
- Redundant servers. If the web site is simply "Here's some info about our company," then it might not be worthwhile to invest in one or more redundant servers. However, if the site contributes directly to your business's revenues, redundancy is critical so that if a web server fails, another can continue to provide service (usually through clustering). In a large web site with frequent traffic, a web server farm probably has several clusters, each comprised of several servers.
- A secure firewall if the site is connected to your LAN. However, it is a better practice to protect normal traffic to and from the LAN and to place web content on a physically separate and independent Internet connection. This is a security precaution ensuring that even if a malicious Internet user breaches your web security, the LAN is still unreachable.

- One or more webmasters to create and manage the content. In a smaller organization, the network administrator might fulfill this function, but larger organizations have full-time employees or contractors.

Most webmasters are experts at creating web content only! They are not web server administrators, and should not be given access to web administration tools.

In the past few years, all major server manufacturers have brought 1U or 2U servers to market. These are dedicated network appliances that serve only one purpose and are not intended for multiple services or applications. High-density, low-cost web servers are increasingly popular for serving web content. The local hard disk only stores the NOS and web server software, while the web content is usually on a separate array, so only a single local hard disk is required. You can also consider using caching-only servers that only cache web content for increased responsiveness. These servers can cache inbound Internet traffic (known as forward proxy) to improve performance for your users or for your web site. Reverse proxy caches your site's content for Internet clients accessing your web servers. This offloads incoming requests for static content, increasing the number of concurrent users or connections the web server is able to maintain, while at the same time improving the browsing experience for those users pointing their browsers to the web server.

After configuring the web servers with adequate hardware, you have another issue to consider: Where do you put the server? You can place the server in your local organization, except that if the site is large and busy, the physical plant might not be sufficient, and building another server room might not be practical. Many organizations co-locate their web servers at a company whose primary business is to provide a highly available physical site for your web. This provides many advantages:

- Co-locators provide the physical plant, including power and environmental controls. Using co-location is generally very cost effective compared to building from scratch.
- Site traffic has no impact on your local network utilization, and because the site is physically disconnected from your LAN, a breach in web security does not directly affect local operations.
- Bandwidth availability is excellent, as most co-locators are directly tapped into the Internet backbone and usually offer extremely high throughput and redundancy through a **SONET** ring, which is a fiber optic transmission medium that is self-healing. If a line is cut, traffic redirects to another ring.
- Co-locators offer guaranteed uptime. While policies vary from one co-locator to the next, redundant power (including UPS systems and backup generators), redundant Internet connections, fire detection and extinguishing, and a high level of security help make sure your site stays available. In addition, co-locators provide on-site personnel 24/7.

- Stringent security requires administrators visiting their servers to have a passkey of some kind to open the door and access the site. Security cameras are everywhere.
- Administrators can still remotely monitor and manage the site over the Internet.

Colloquially speaking, administrators refer to a co-location arrangement as "ping power pipe." You can "ping" the server from anywhere on the Internet to see if it responds, "power" is always available, and the "pipe" is the line to the Internet.

In Hands-on Project 9-4 at the end of this chapter, you will visit the web site of a co-locator.

Configuring a Web Server

Each major NOS offers a different product for web management, as listed below:

- OS/2 and NetWare 5.1—WebSphere: an IBM product that provides excellent management tools with a relatively intuitive graphical user interface. OS/2 can also use a combination of WebSphere, HTTP Server, Lotus Notes Domino, and Apache.
- UNIX/Linux—Apache: a text-based tool that (as usual for UNIX products), is open source code, highly reliable, and not at all intuitive. However, a newer version of Apache for Windows NT is in development and should be more intuitive.

Perhaps the reliability of UNIX and Apache is best attested to by Microsoft itself. For years, Microsoft ran a version of UNIX known as FreeBSD and Apache web services for the Microsoft Hotmail web site. In 2000, Microsoft finally switched to their own products: Windows 2000 and IIS 5.0.

- Windows NT/2000—Internet Information Server (IIS): the best testament to IIS is that Microsoft's own web site, one of the largest and busiest in the world, runs IIS 5.0. The GUI is simple to operate, yet is rich in features.

Each of these products is a study in itself; however, Hands-on Project 9-4 at the end of the chapter will step you through the basics using IIS 5.0 on Windows 2000 Server. In addition, there are scores of third-party web server software products vying for position within the ever-expanding web server market.

Remote Access Service

Remote access service (RAS) is the ability of a server to accept a connection from a client even when physically disconnected from the LAN. Users establish the connection on their end through dial-up modem connections or existing Internet connections using a **virtual private network (VPN)**. Once the connection is established, the user

experience is the same as if directly connected to the local LAN, except that over a dial-up connection, network responsiveness is much slower. A VPN connection can be much better because although it usually takes place over the public, unsecured Internet, the session is protected inside an encrypted virtual "tunnel" that is extremely difficult for intruders to breach. The VPN can be just as slow (or slower because of the encryption overhead) if the Internet connection is dial-up, but with the availability of high-speed Internet access, a user might experience significantly better performance.

Regardless of the method, RAS represents another arena of responsibility for the administrator. You provision the server to accept RAS connections, and you also ensure that the connections are secure. The physical preparations for a modem pool were discussed at the end of Chapter 7. In Hands-on Project 9-7, you will use Windows 2000 as an example of configuring a server so that it is ready to accept the connections.

You can also create VPNs inside your LAN to further guard against intruders. For example, if you wanted to protect a server that contained sensitive information, you could open it up only to users permitted to connect through a VPN (in addition to setting appropriate file and logon permissions as usual). Anyone intercepting the VPN session between user and server is very unlikely to retrieve any useful information.

Although a persistent hacker could obtain the phone numbers for dial-up connections, you should make it as difficult as possible. Instruct users not to share the phone numbers with persons outside your organization. If you suspect a compromised number, change it as soon as possible.

Protocol Support

You are already aware of the basic protocols such as TCP/IP, IPX/SPX, and NetBEUI. You can usually use each of these protocols to communicate with a RAS server. However, there are additional protocols that are used to establish and secure the connection.

The first type of protocol is a line protocol. Network protocols are designed for network media such as 10BaseT Ethernet over CAT5 cable. When RAS clients dial in to the server over a phone line, TCP/IP alone won't work, because it is incompatible with conventional phone lines. Instead, when clients dial up the modem and establish a connection, a line protocol encapsulates the network protocol. Encapsulated packets are then sent across the connection where the server unwraps the line protocol packet. The packet then transmits over the network like an ordinary network packet. This process is illustrated in Figure 9-12, in which the PPP packet encapsulates a TCP/IP packet.

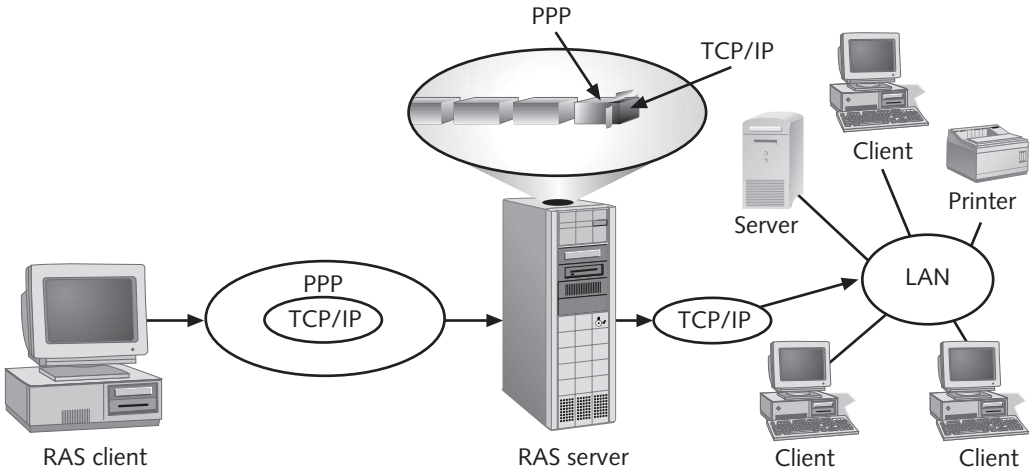

Figure 9-12 Line protocol encapsulation

The following are the two most common line protocols:

- **Point-to-Point Protocol (PPP)**. A very flexible line protocol, PPP interoperates with a variety of RAS software packages. PPP supports the NetBEUI, IPX/SPX, and TCP/IP protocols, data compression and encryption, and authentication protocols (addressed later in this chapter). The NetWare implementation of PPP is PPPRNS (Point-to-Point Protocol Remote Node Service) and is very similar, though it does not support NetBEUI.
- **Serial Line Internet Protocol (SLIP)**. A more common line protocol in the past than it is today, SLIP uses only the TCP/IP protocol and is useful for UNIX connections. In Windows, when you connect with SLIP, a Windows Terminal dialog box opens allowing you to perform an interactive logon with the UNIX server. SLIP is very basic and does not support authentication protocols, encryption, or compression.

In addition to a line protocol, you need a tunneling protocol in order to establish a VPN connection. There are two primary tunneling protocols:

- **Point-to-Point Tunneling Protocol (PPTP)**. A popular and easy-to-configure tunneling protocol. Configure both the server and the client to establish the VPN connection using PPTP and make the connection.
- **Level 2 Tunneling Protocol (L2TP)**. A newer VPN protocol that requires an established certificate authority. Clients establishing a connection must download a certificate from the certificate authority. The certificate then validates the connection attempt over the VPN connection attempt.

Security Protocols

Logging on remotely exposes the connection to eavesdroppers who could tap the connection and retrieve user name/password combinations. To avoid this vulnerability, you need to select an encryption method. The most common ones are listed below and are supported in varying degrees by all the major NOSs.

- **Password Authentication Protocol (PAP)**. PAP sends logon information in clear text—using a packet sniffer, an eavesdropper can analyze the packet and retrieve the logon data. PAP is the last resort—use it only when the server you dial into does not support any of the other authentication protocols.
- **Shiva Password Authentication Protocol (SPAP)**. Shiva products (acquired by Intel) are a popular alternative to Microsoft RAS solutions. Shiva encrypts authentication credentials for Shiva LAN Rover software.
- **Challenge Handshake Authentication Protocol (CHAP)**. CHAP is a flexible and common authentication protocol that supports encryption for a variety of operating systems. Microsoft has two specific implementations: MS-CHAP for all Windows clients and MS-CHAP v2 for Windows 2000 clients.

Configuring RAS on a Server

Each network operating system has unique methods for configuring a RAS server. The important thing is to know about all the protocols discussed so far in this chapter, because a critical matter in configuring client/server RAS communication is ensuring that settings on both sides match. Mismatched settings do not cause a poor connection; they usually prevent connection altogether. You must then determine what caused the problem. Most of the time, connection problems relate to mismatched security protocols. For example, if you are trying to connect using PAP but the RAS server accepts only MS-CHAP v2, the client logon attempt will be denied.

Figure 9-13 shows a Windows 2000 dial-up client with authentication settings for PAP, SPAP, and CHAP. As a rule, connection attempts first try the most secure available method and, failing that, move to the next most secure method until all available authentication methods are exhausted. Figure 9-14 shows the default Windows 2000 Server dial-in settings. The client shown in Figure 9-13 will not succeed in the connection attempt to this server, because only MS-CHAP and MS-CHAP v2 are permitted.

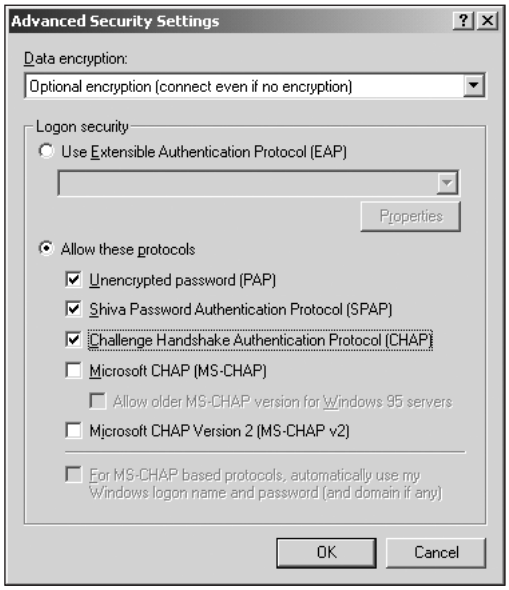

Figure 9-13 Windows 2000 dial-up client authentication

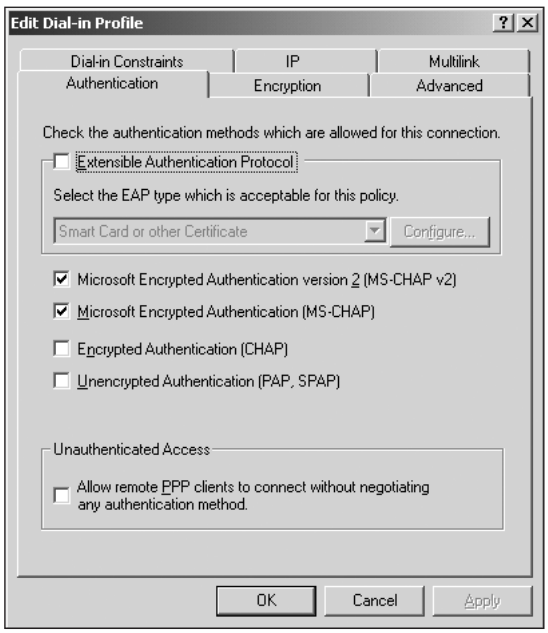

Figure 9-14 Windows 2000 Server dial-in settings

Fax Services

Faxing has become a vital business activity. While some organizations use fax machines and only require light usage, others might send hundreds or thousands of pages each day. For organizations that require high-volume fax capability, it becomes completely impractical to stand in line at the fax machine hoping that someone ahead of you is not faxing *War and Peace*. Fortunately, a combination of hardware and software technology makes sending and receiving faxes much more efficient and productive, and users can perform fax functions from their desks.

Even with the best fax server systems, most organizations will still keep some fax machines around, as it is still practical for sending copies of physical documents. Also, if you want to send a quick fax to someone and do not expect to do so again, it's probably easier to just punch in the number and hit "send" than to enter all the contact information in the fax software.

A fax server comes in two basic hardware formats: turnkey and computer-based. **Turnkey fax servers** are self-contained, freestanding devices in which fax software and hardware are already installed, and except for some company-specific configurations, they are ready to use right out of the box. Our attention will focus on the alternative, computer-based fax servers, which require separate hardware and software.

Fax hardware can consist of a simple fax modem that is shared from a single PC. While this might be acceptable for a small workgroup requiring only light-duty fax service, our focus will be on higher-volume solutions for which you will need one or more fax boards. Besides the fact that a single fax board can send or receive many times more faxes than a conventional fax machine, a fax board offers several other characteristics and advantages:

- Fax boards are specially tuned for fast **handshakes** (that soothing squawking noise that faxes and modems make when establishing a connection) and fax compression. This only saves a few seconds, but when multiplied by the number of users and the long distance rate of faxes (especially overseas), it makes a big difference. In addition, fax boards have a higher throughput than conventional faxes, so even at 14.4 Kbps, it can send a page in about 15 seconds instead of 60 seconds.

The most common fax speed remains 14.4 Kbps even though there is also a 28.8 Kbps standard available.

- Resolution is usually better on a fax server than most conventional fax machines. Also, the fax software often includes the ability to "clean up" unclear faxes.

- Security-sensitive documents are vulnerable to the curious on a conventional fax machine. Once at a company where I worked, a manager faxed a spreadsheet of all the salaries for employees in my department and left it on the machine. A co-worker found and read the spreadsheet. (No matter what I did, I couldn't get him to tell me other people's salaries; perhaps he knew it would make me mad.) On a fax server, you fax from the desktop and nobody else sees it.
- A fax server saves users the frustration of clearing jams, replacing toner and paper, waiting in line, and beating the machine with other office equipment.
- Most fax machines collect data on inbound and outbound faxes, but you print it out and that's that. A fax server allows you to organize the data and use software products specially designed for reporting (such as Crystal Reports). Reports are useful for determining which users or departmentrs send and receive the most faxes, finding out when fax traffic is heaviest, and so on. Reports are also as permanent a record as you like, whereas fax machine logs must be periodically cleared.
- Inbound faxes can go to a departmental printer, or remain electronic and go to fax software or, better yet, an email inbox on the user's computer.
- The fax software automatically fills in cover sheet information, such as who is sending the document and the number of pages. Broadcast fax cover sheets can name individuals instead of all parties.
- Mobile users might be unavailable to receive a physical fax at the office. Using client fax software, they can have the fax forwarded to their email inbox or to a printer at their current location.

Fax boards are PCI, and therefore can be installed in most computers—PC or otherwise. However, fax server software is usually a UNIX, NetWare, or NT/2000 product. Many vendors also include software that integrates with Lotus Notes, Microsoft Exchange, Novell GroupWise, SMTP/POP3, and more.

You can combine other services or functions with a fax server, but be aware of the kinds of resources it requires. First, faxing requires processing power for compression, error correction, data packet creation, and other fax-related processes; so you should not share any processor-intensive functions with the fax server. An exception for this can be made if the fax board itself offloads these duties from the processor. Second, if faxes are stored on the server, plan for plenty of hard disk space. Calculate for the average size of inbound faxes and the length of time they are likely to be stored. In many cases though, you will offload the faxes to another server such as a Novell GroupWise or Microsoft Exchange Server. Memory requirements are not high; you usually only need enough to run the fax software, probably 32 MB.

For more about fax servers, consult:

- *www.faxserverfaq.com*
- *www.dialogic.com*
- *www.facsys.com*

SNA Server

Systems Network Architecture (SNA) Server harkens back to 1974 and IBM's networking standards for mainframes. In those days, users sat at dumb terminals that accessed a mainframe server through a text-based terminal session. The reason why this history is important is that the basic mainframe architecture still exists today, and access to many midrange systems is similar—IBM's AS/400 and AS/390 servers, for example. Users must connect to these servers, but usually do not sit at dumb terminals anymore. Now, users usually sit at Microsoft Windows clients and connect to the mainframe using a terminal emulator that appears to the server as if it were a dumb terminal. However, making a straight connection like this from Windows and maintaining all the features of Windows and the mainframe at the same time is not always practical or possible. SNA Server is a Microsoft product that acts as a gateway between the client and the mainframe.

Microsoft has a successor to SNA Server: Host Integration Server, which is compatible with Windows NT 4.0 and integrates with Windows 2000. For the sake of our discussion, we will collectively refer to these products as SNA Server.

The following list briefly describes several main features of SNA Server:

- Windows users connect to the mainframe (known in this context as a "host," not to be confused with simply a TCP/IP host) through client software. From a web browser, a user can click on a link that automatically downloads the necessary client software to their local computer. The client software allows users to be a "client" to the host and run whichever terminal emulation software or application suits their needs.
- Users have difficulty remembering passwords, and experience even more difficulty if they must log on more than once—first to a Windows 2000 server and then to a host. Often, the accounts are not synchronized (the user name and/or password is different), further complicating matters. Using SNA Server, administrators can synchronize user accounts so that they are consistent between the host and the Windows NT or Windows 2000 domain, and changing the password on one system automatically synchronizes with the other.
- Database availability is improved because SNA allows users to access mainframe-based databases such as DB2 (an IBM database).
- Uniform performance monitoring allows the administrator to use Windows NT/2000 Performance Monitor to analyze server performance on hosts.
- Administrators streamline the burden of file resource access. SNA applies Windows NT/2000 file security on shared folders as if the resource was on a local Windows NT/2000 server. This allows you to apply the same security permissions and access rights as with any other file.

- Print compatibility is improved because users can send mainframe print jobs to LAN printers without changing host applications.
- SNA server failure has less impact on user productivity because you can use multiple SNA servers for failover and load-balancing purposes.
- By itself, a logon using a 3270 terminal, for example, is sent in clear text, making it easy for someone intercepting the sign-on to read the user name and password. SNA Server encrypts the data streams between the server and client.

Application Servers

You can run applications in one of three basic implementations: dedicated, distributed, or peer-to-peer.

- **Dedicated application**—As discussed with some of the services earlier in this chapter (email services, for example), servers sometimes run only a single application or service and nothing else. This helps to assure that application performance is unhindered by interference from other applications or services and also contributes to the stability of the server. The application software does not run on the workstation, though a client piece might, such as using Outlook Express to retrieve email from a UNIX Sendmail dedicated server.
- **Distributed application**—The application runs on the server. The client can send requests to the server but does not run the application or perform processing. Probably the most common example of this type of application is a database. The client sends a query to the database server, which returns the results of the query to the client.
- **Peer-to-peer (P2P) application**—A good example of a peer-to-peer application is the controversial Napster music-sharing service. A P2P application server primarily exists to run software that allows peer computers to communicate with one another. In the case of Napster, the peers search for and download music from one another. The server functions mostly as a gateway for the clients. An important issue with P2P servers is that they provide appropriate security measures so that users can search for and download files from other users without being able to access unauthorized materials.

Monitoring Protocols

Most networks are much larger than the administrator can practically administer if he or she had to physically visit each server and network device. Therefore, a means by which administrators can access remote servers is necessary. Each network operating system allows general administration from afar using various administrative tools. For example, you can manage a NetWare 5.1 server using the NWAdmn32.exe program from any Windows client regardless of geographic distance. However, administrators also require

a way to proactively monitor servers and network equipment like routers, hubs, and switches. How does the administrator know when a router becomes saturated beyond its ability to function? Is there a way that administrators can see an increasing trend in network traffic and prepare a proactive solution? The answer to both questions is found in the monitoring protocol.

There are two primary monitoring protocols: Simple Network Management Protocol (SNMP) and Desktop Management Interface (DMI), with SNMP being the most prevalent.

SNMP

Simple Network Management Protocol (SNMP) is really of no use by itself. Its usefulness is comprised of several elements that work together with the ultimate purpose of informing the administrator of a changing trend in the use of an object or alerting the administrator of an error, failure, or condition. For example, an administrator might want to know well in advance when the 100 GB disk array is down to 10 GB so that he or she can make sure that more storage is available. This prevents the reactive response in which users complain about out-of-disk-space errors and the administrator must scramble at the last minute to find additional storage. With third-party SNMP software, the administrator can be notified in a number of ways about the low-storage problem. The administrator can configure a response in any of several forms, such as an email message, pager alert, or network message to a workstation at which the administrator is logged on.

A useful SNMP solution consists of several individual components:

- **SNMP management system**—Also known as a management console, this is the computer that runs SNMP management software. The software can run on any compatible computer; it does not have to be a server, though it often is. The SNMP management system sends requests for information from the monitored system, known as the SNMP agent. Except for an alarm-triggering event, the SNMP agent does not normally initiate messages.
- **SNMP agent**—A service that runs on the actual object you want to manage or monitor. For example, if you have a web server and want to receive an alert when the number of concurrent connections exceeds a certain threshold, you would configure the web server as an SNMP agent.
- **Management Information Base (MIB)**—A database of definitions for the specific device being monitored. For example, one of the MIB values associated with a particular device could be "sysUpTime," which specifies the elapsed time since the managed device was booted. Similar to host-name-to-IP-address mapping, a MIB value has an official name and a dot notation. For example, the dot notation for sysUpTime is 1.3.6.1.2.1.1.3.0, though it's obviously easier to use the official name.

- **Communities**—A group of hosts that use the same community name. You can name the community whatever makes sense for your organization. The community name is not so much a grouping as it is a small measure of security. When SNMP queries are issued to a community, only members of that community respond, and the community name functions as a rudimentary password.

SNMP has no security of its own, and a malicious user with a packet sniffer could retrieve SNMP traffic, and then through impersonation, return false information to the SNMP management system. Newer implementations of SNMP are introducing public key/private key verification techniques to prevent this situation.

- **Traps**—When the SNMP agent issues a message to the SNMP management system, it is known as a trap.

As an example of configuring an SNMP agent, look at the Windows 2000 SNMP Service Properties dialog box in Figure 9-15. The check boxes in the Service frame specify what type of device is involved. For example, if the device is a router, you would select the Internet item. For explanation of the other devices, use the online Help function.

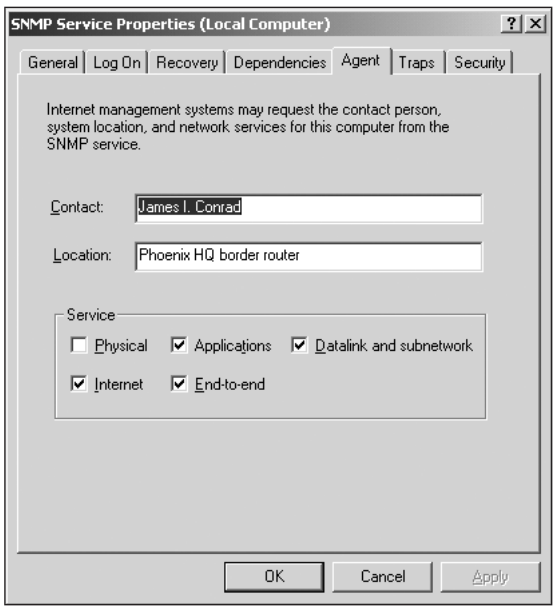

Figure 9-15 Configuring the SNMP agent properties on a Windows 2000 server

Continuing the Windows 2000 example, notice the Traps tab shown in Figure 9-16. Here, you specify the community name and the trap destination. The agent can only send messages to hosts that know the community name.

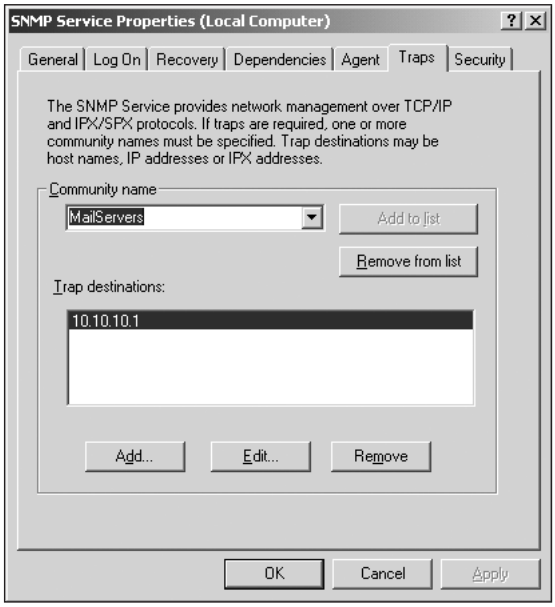

Figure 9-16 Specify the community name and the trap destination's IP or IPX address

Finally, specify rudimentary security on the Security tab, where you select accepted management systems with which the agent is allowed to communicate (Figure 9-17).

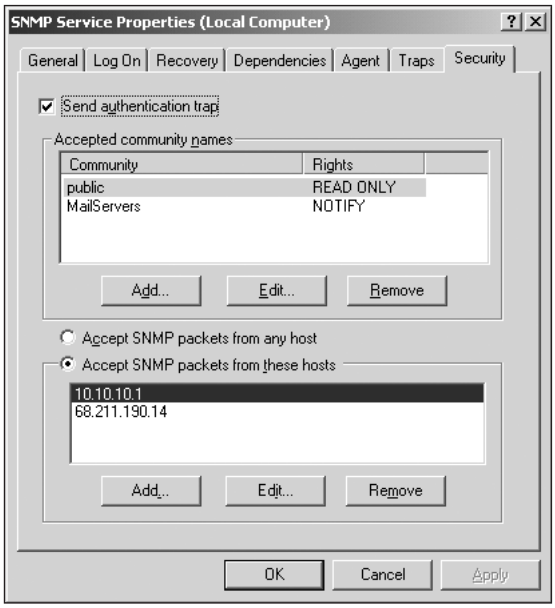

Figure 9-17 Use the Security tab to specify SNMP management systems with which this host can communicate

You might also find RMON on certain network devices. **Remote monitoring (RMON)** is an extension of the SNMP protocol and provides more comprehensive network monitoring capabilities. Instead of devices answering queries from the SNMP management system, RMON proactively sets off alarms for a variety of traffic conditions. As the full RMON protocol is quite comprehensive, only portions of it are usually placed into network devices such as routers.

There are several utilities and applications that use SNMP; however, the three most significant vendors are:

- Computer Associates' Unicenter TNG (*www.ca.com*)
- IBM's Tivoli (*www.tivoli.com*)
- Hewlett-Packard's OpenView (*www.hp.com*)

DMI

The **Desktop Management Interface (DMI)** is similar to SNMP, except that it contains specific information about an actual device. Instead of using a MIB, DMI uses a **Management Information File (MIF)** database that can contain information such as model ID, serial number, memory, and port addresses. DMI often runs in conjunction with SNMP. For example, when an SNMP query arrives at the agent, DMI can enter MIF information in the SNMP MIB.

The Server as a Network Device

Network operating systems can enable a server to fulfill roles traditionally reserved only for hardware network devices such as routers, bridges, and firewalls. I recommend that you stick with the faster and more reliable hardware solutions for the most part, because servers must simultaneously cope with countless other activities and variables. For example, how often does a router crash because of a software bug? Never. But a server functioning as a router must not only perform routing functions, but also simultaneously load millions of lines of operating system code into memory, run a few services just to stay alive (if not perform additional functions), share processor utilization with other functions, depend upon many other physical devices in the server (each one representing another possible point of failure), and so on.

Nevertheless, for the smaller, value-conscious organization or a smaller department within an organization, a server network device might suit your needs. For example, even if you work for a gigantic organization, if it has a small satellite office in Nome, AK, with only five workstations and one server, it's probably not cost effective to install a Cisco router costing thousands of dollars to connect the office to a WAN. Instead, you would use the existing server. Because there is likely to be very little network traffic going through the server anyway, the risk is reduced to the point of being acceptable.

The Server as a Router

Configuring the server as a network device such as a router (also known as a gateway for our purposes) requires at least two network cards. In the case of a server configured as a router, one NIC connects to one network or subnet, and the other NIC connects to another network or subnet. The static IP address you assign to each respective NIC is in the same network range as the subnet to which it connects, as illustrated in the NetWare 5.1 example Figure 9-18. You will also configure the NICs the same way if the server functions as a firewall or bridge.

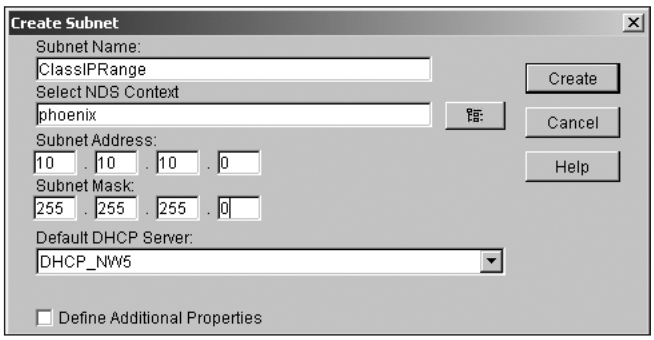

Figure 9-18 The NIC for each respective network matches the network address of the subnet to which it connects

A server router connecting two different types of networks (for example, Ethernet and Token Ring) also acts as a bridge. There is not a specific configuration in most NOSs to make the server a bridge only.

After configuring the NICs, it's only a matter of configuring the respective operating system to perform the function you want. For example, to configure Windows 2000 as a router, you would use the Routing and Remote Access (RRAS) tool (Start, Programs, Administrative Tools, Routing and Remote Access). If this is the first RRAS function you have configured, a wizard guides you through the process. Otherwise, you will have to manually add a routing protocol such as RIP or OSPF and then specify the NICs participating in the routing function. (A brief explanation of these routing protocols follows.)

RIP

Router Information Protocol (RIP) is a distance vector-based protocol, which identifies the best route for a destination based on the number of hops (number of routers that the message passes through). RIP has the following advantages:

- *Configuration*—You enable RIP on the router interface; no further configuration is necessary.

- *Simplicity*—Routing table advertisements are very simple (compared to OSPF) and easier to read.
- *Compatibility*—Most routers are compatible with RIP.

RIP is a fine protocol for simple purposes, but it can run into problems in more complex routing scenarios. Its disadvantages include:

- *Delayed convergence*—**Convergence** is the state in which all routers connected to a network have the same updated information. RIP exchanges routing tables in an unsynchronized and unacknowledged manner. Consequently, internetwork convergence might take several minutes, a costly lapse in a production network.
- *Larger routing table size*—Multiple routes to a network destination can appear as multiple entries in the routing table. Though this might not affect smaller networks, large networks with multiple paths can generate hundreds or thousands of entries in the RIP routing table.
- *Greater network traffic overhead*—Because of the larger routing tables generated by RIP, route advertising requires more network overhead. Route advertising continues to update routing information every 30 to 90 seconds, even after convergence. This means that the router's entire routing table is passed to all other known routers.
- *Poor scalability*—RIP broadcasts its entire routing table, making them much larger than OSPF tables. This generates an unreasonable degree of network traffic on larger networks. Also, RIP only accepts up to 15 hops.
- *Routing loops*—When a router learns of a downed link, it notifies a neighbor, which in turn notifies its neighbors. However, if an upstream neighbor issues its 30–90 second update table before it learns of the downed link, then all its downstream neighbors view the downed link as available again. This routing loop is difficult to troubleshoot and usually requires manual intervention to fix.

The routing loop description here is only one type of routing loop and is known as "counting to infinity." Several other types of routing loops could occur within a RIP network.

- *Less efficient*—Because RIP builds best path determination based only upon the number of hops to a destination, the path can sometimes be less efficient. For example, it might take only three hops to get from Network A to Network F. However, one of the hops must take place over a 56 Kbps demand-dial link. A faster route might exist over a four-hop T1 link, but a RIP-based router will not use it.

OSPF

The **Open Shortest Path First (OSPF)** routing protocol builds its routing tables using a link-state algorithm, which calculates the shortest path to each host based on the network topology, not just the fewest number of hops (as with RIP). The algorithm uses not only hop count, but also other factors, such as available bandwidth and network congestion, to determine the best path to a destination. Routing tables are smaller and update more often, but they are more efficient and converge more quickly than RIP. OSPF is quickly replacing RIP in most implementations. OSPF has the following advantages:

- *More rapid convergence*—OSPF routers converge much more quickly than RIP routers and are not susceptible to routing loops.
- *Smaller routing table size*—Instead of storing every possible path to a destination, the SPF algorithm calculates the best path and stores only that path in the routing table.
- *Less network traffic overhead*—Because of smaller routing tables, and because routing tables are not broadcast redundantly, OSPF requires less network traffic.
- *Hop count*—There's virtually no limit on hop count.
- *Scalability*—All of the above features make OSPF more scalable and suitable for large networks.
- *Compatibility*—Most current network devices can accept and use the OSPF protocol.

The OSPF routing protocol also has some disadvantages, but they are largely outweighed by the advantages:

- *Greater use of resources*—The database of link-state advertisements and the SPF calculations required for path determination require more memory and processor use than RIP.
- *More complex configuration and use*—Link-state protocol implementation requires careful planning, and the configuration options are more extensive than for RIP. Link-state protocols are also more complex and difficult to understand than RIP; however, this might not be a direct concern unless you need to perform detailed network analysis.

As a rule, most network administrators prefer the scalability, flexibility, and accuracy of the OSPF protocol over RIP.

The Server as a Firewall

Recall from earlier discussions that a firewall protects your LAN from other networks (usually the Internet). A proxy server serves the same function by impersonating the internal client, and is a term we will use interchangeably with firewall. The proxy accepts

Internet requests from the LAN client, represents the client on the Internet, and issues the request to the Internet destination host. When the Internet host returns an answer (usually a web page), the proxy forwards the content to the original LAN client. The proxy server performs this function for all its internal clients and, to the Internet at large, appears to be a single, very busy Internet client. The proxy also prevents LAN users from accessing sites forbidden by the organization.

A firewall server can be a hardware or software solution, and is in itself a form of routing. However, server firewalls can offer very sophisticated and advanced features, making potential weaknesses in operating stability a more acceptable risk, especially when there are redundant proxy servers.

Besides protecting the LAN, a proxy server improves Internet performance for the clients by using caching. When the proxy server delivers a web page to a client, it also caches the page in memory and on a hard disk cache. The proxy server delivers the same content to the next client that wants to view the same page. This saves Internet network bandwidth and greatly improves responsiveness for frequently accessed pages.

Although you can configure the basic NOS to perform some firewall functions, such as allowing access only to certain IP addresses, you configure a proxy server using additional server software in most cases. For example, Novell NetWare uses an add-on product called BorderManager, and Windows NT/2000 use Microsoft Proxy Server 2.0 or Internet Security and Acceleration (ISA) Server.

Chapter Summary

- DHCP automatically distributes IP addresses to DHCP clients, avoiding the error-prone and time-consuming process of manually entering IP configuration on hosts. The DHCP lease process uses a discovery, offer, request, and acknowledgment process.
- DNS resolves IP addresses to names and vice versa, replacing the need for a HOSTS file. The resolver is the client waiting for a name resolution answer. If a DNS server does not have the answer to a request, it can use other DNS servers in a recursive query or use an iterative query, which simply refers the original resolver to another DNS server.
- DNS servers are useful for forward lookups, in which a resolver submits the host name and requests an IP address, or reverse lookups, in which the client submits the IP address and requests the host name. There are several types of DNS records, including A records, MX records, and CNAME records.
- There are several types of DNS servers, including primary master, master, slave, secondary domain, and caching-only.
- Beginning with NetWare 5.1 and Windows 2000, you can rid the network of reliance on NetBIOS because you can instead use Dynamic DNS (DDNS). Because of an interaction with the DHCP service, DDNS can accept automatic registrations

from clients when they receive their IP configuration from the DHCP server. Any environment including UNIX that implements RFC 2136 can use DDNS.

- WINS is a method by which user-friendly NetBIOS computer names can be resolved to IP addresses for the purpose of allowing such communication between hosts on different subnets. Name registration occurs when a WINS client requests the use of a NetBIOS name from the WINS server. The WINS server can either accept or reject the request for a NetBIOS name made by the WINS client. The response that is given depends on several factors.
- Name resolution on a WINS network attempts name resolution in the following order: DNS (if the name is more than 15 characters or contains periods), name cache, WINS server, broadcast, LMHOSTS file, HOSTS file, and DNS. As a memory aid, you can take the first letter of each step in the order and correlate it to the first letter of the following sentence: "Can We Buy Large Hard Drives" (Can = Cache, We = WINS, Buy = Broadcast, and so on).
- No matter how many WINS servers you have in your network, you will still only have a single WINS database. With multiple servers and one database, replication is extremely important to ensure that all of the WINS servers have a consistent copy of the database.
- WINS replication occurs through pull replication partners, which request WINS updates from their configured push replication partners. Only the changed or new WINS records replicate, not the entire WINS database. If a client is not WINS-compatible, a WINS proxy agent can register on behalf of the non-WINS client.
- Mail services are a critical part of most organizations. Email requires one or more of several protocols, including SMTP, MIME, POP3, and IMAP. Main email server applications include Lotus Notes, Novell GroupWise, Microsoft Exchange, and UNIX/Linux Sendmail. Whatever your choice, you should carefully protect your email with redundant mail servers for failover purposes. Also make sure additional DNS servers are available to direct mail to the appropriate domains.
- When configuring your mail server, be sure to open the firewall to mail-related traffic, add appropriate MX records in DNS, and supply plenty of RAM and hard disk space. Be sure to defragment the mail server hard disk frequently. For best performance and reliability, dedicate the server to only mail. Have a comprehensive anti-virus solution to protect the network, which includes the servers and especially all workstations with email clients because mail is a common carrier of viruses.
- Your web servers should have plenty of Internet bandwidth, which is best accomplished by separating it from LAN bandwidth using a dedicated Internet connection. Plan for plenty of fast hard disk space for file downloads and streaming multimedia. Users will trust your content more (especially online transactions) if you have a digital certificate.
- Install enough memory so that much of the web content can be cached, and install redundant web servers for failover. Protect your LAN from attacks on the web servers by separating it with a firewall. Do not allow webmasters to administer the web site itself.

- A co-location site can provide "ping power pipe" availability for your web with excellent bandwidth availability and redundancy, redundant power, security, and regulated environmental controls.
- Common web server software includes Apache, WebSphere, Lotus Notes Domino, and Microsoft IIS.
- Remote access service (RAS) is the ability of a server to accept a connection from a client even when physically disconnected from the LAN. Users establish the connection on their end through dial-up modem connections or existing Internet connections using a virtual private network (VPN). Once the connection is established, the user experience is the same as if directly connected to the local LAN, except that over a dial-up connection, network responsiveness is much slower.
- RAS uses basic network protocols, but also requires a line protocol such as PPP or SLIP; security protocols such as PAP, SPAP, CHAP, MS-CHAP, or MS-CHAP v2; and tunneling protocols such as PPTP or L2TP. One of the most important things to remember about RAS server configuration is that the RAS client settings must match or else the connection will fail.
- Fax servers offer advantages over fax machines such as greater speed, improved resolution, better security, better reporting, routing, and forwarding to another fax machine or email inbox. Fax servers consist of a fax board and fax software.
- SNA Server acts as a gateway between network clients and mainframe or midrange hosts. Clients then connect using terminal emulation. Users can gain access to large mainframe databases such as DB2, administrators can monitor performance on hosts, host print jobs can go to LAN printers, and logon security is improved.
- SNMP is a management protocol. Its usefulness is comprised of several elements that work together with the ultimate purpose of informing the administrator of a changing trend in the use of an object or alerting the administrator of an error, failure, or condition. SNMP functionality relies on an SNMP management system, an SNMP agent, a MIB, traps, and communities.
- Desktop Management Interface (DMI) is similar to SNMP, except that it contains specific information about an actual device. Instead of using a Management Information Base (MIB), DMI uses a Management Information File (MIF) database that can contain information such as model ID, serial number, memory, and port addresses. DMI often runs in conjunction with SNMP.
- A server can function as a network device such as a router (using RIP or OSPF) or a firewall/proxy. Routing functions are best performed by hardware routers unless it is not necessary or cost effective.
- RIP is simple but limited to 15 hops and susceptible to routing loops. OSPF is practically unlimited in the number of hops and uses a routing algorithm for creating routing tables. OSPF is much more efficient than RIP.

- The server can also act as a firewall/proxy server to protect the LAN from malicious Internet users and prevent LAN users from accessing forbidden web content. Besides this protection, a proxy can cache web content for faster delivery of web content and reduced burden on Internet lines.

KEY TERMS

acknowledgment — In DHCP, a confirmation from the server to the client that the DHCP lease process completed successfully.

address (A) record — Also known as a host record, this is the actual record that resolves the host name to the IP address.

BOOTP — The Bootstrap Protocol; uses a BOOTP server that can distribute IP addresses to clients (similar to DHCP).

caching-only server — A DNS server that has no zone database, either of its own or copied through a zone transfer from a primary domain server. Caching-only servers mostly function to improve performance by reducing the number of forwarded queries.

Challenge Handshake Authentication Protocol (CHAP) — A flexible and common authentication protocol that supports encryption for a variety of operating systems. Microsoft also has two specific implementations: MS-CHAP for all Windows clients and MS-CHAP v2 for Windows 2000 clients.

CNAME record — Stands for a canonical name record and is an alias that points to another host.

communities — A group of hosts that each use the same community name. You can name the community whatever makes sense for your organization. The community name is not so much a grouping as it is a small measure of security. When SNMP queries are issued to a community, only members of that community respond, and the community name functions as a rudimentary password.

convergence — A state in which all routers connected to a network have the same updated information.

daemon — The Linux and UNIX name for a service.

dedicated application — A server running a single application or service and nothing else. This helps to assure that application performance is unhindered by interference from other applications or services and also contributes to the stability of the server.

Desktop Management Interface (DMI) — Similar to SNMP, except that it contains specific information about an actual device.

DHCP server — A server that automatically allocates IP address configuration to DHCP clients.

discovery broadcast — A broadcast initiated by a DHCP client that seeks a DHCP server.

distributed application — The application runs on the server. The client can send requests to the server but does not run the application or perform processing.

DNS zone — A naming boundary for which a DNS server is responsible.

Domain Name System (DNS) — A service that stores a record of both the node's IP address and host name, and uses these records to service name resolution requests.

Dynamic DNS (DDNS) — Automatic registrations from clients at the time they receive their IP configuration from the DHCP server.

Dynamic Host Configuration Protocol (DHCP) — A protocol that allows its clients to lease IP address configuration automatically from a DHCP server.

handshake — The squawking noise that faxes and modems make when establishing a connection.

HOSTS file — A plain text file that contains static, manual entries of host-to-IP-address mappings.

Internet Message Access Protocol (IMAP) — In its current version (IMAP4), IMAP allows the email client to leave messages on the mail server even after logging on instead of downloading each one.

iterative query — When a DNS server refers the resolver to another DNS server that might be able to resolve the request.

lease — The length of time for which a client receives IP configuration from a DHCP server.

Level 2 Tunneling Protocol (L2TP) — A relatively new VPN protocol that requires an established certificate authority. Clients establishing a connection must download a digital certificate from the certificate authority. The certificate then validates the connection attempt over the VPN connection attempt.

LMHOSTS file — A plain text file that contains static, manual entries of NetBIOS name records.

Mail Exchanger (MX) record — Routes mail to the appropriate server(s) for members of the domain.

Management Information Base (MIB) — A database of definitions for the specific SNMP device being monitored.

Management Information File (MIF) — A DMI database of information such as model ID, serial number, memory, and port addresses.

master server — An authoritative DNS server that transfers zone data to one or more slave servers. ("Authoritative" means that the server is configured to host the zone and return query results.)

Multipurpose Internet Mail Extensions (MIME) — A protocol that adds the mail capability of attaching and transferring multimedia file attachments. To use MIME, you must also have an email client capable of decoding the MIME format.

Name Server (NS) record — Specifies what DNS servers are delegated servers for the domain, meaning that the server specified in the record can resolve queries authoritatively.

NetBIOS — Broadcast-based name resolution scheme where a client simply broadcasts the NetBIOS name of the computer it wishes to reach to all of the computers on a subnet. The broadcast message identifies a computer that acknowledges the broadcast and establishes a communication link.

offer — The DHCP server's proposed IP address and configuration to the DHCP client.

Open Shortest Path First (OSPF) — A routing protocol that builds its routing tables using a link-state algorithm, which calculates the shortest path to each host based on the network topology, not just the fewest number of hops (as with RIP).

Password Authentication Protocol (PAP) — A security protocol that sends logon information in clear text. Using a packet sniffer, an eavesdropper can analyze the packet and retrieve the logon data.

peer-to-peer (P2P) application — The server primarily exists to run software that allows peer computers to communicate with one another.

Point-to-Point Protocol (PPP) — A very flexible line protocol that interoperates with a variety of RAS software packages. PPP supports the NetBEUI, IPX/SPX, and TCP/IP protocols, data compression and encryption, and authentication protocols.

Point-to-Point Tunneling Protocol (PPTP) — A popular and easy-to-configure VPN tunneling protocol.

Post Office Protocol 3 (POP3) — A line protocol that allows messages to be stored on the mail server for incoming email.

primary domain server — The starting point of all DNS records. The zone database is readable and writeable on the primary domain server: You can add, remove, or modify DNS records.

primary master server — There is only one primary master server per DNS zone, and it is the first and final authority for all hosts in their domain. Primary masters are the source for records that are copied to master or slave DNS servers.

PTR record — The actual record used in reverse lookups.

pull replication partner — In WINS, a replication partner that requests and then accepts changes from its push replication partners.

push replication partner — In WINS, a replication partner that responds to requests for changes from its pull replication partners.

recursive query — A query forwarded from one DNS server to another.

remote access service (RAS) — The ability of a server to accept a connection from a client even when physically disconnected from the LAN.

remote monitoring (RMON) — An extension of the SNMP protocol, providing more comprehensive network monitoring capabilities. Instead of devices answering queries from the SNMP management system, RMON proactively sets off alarms for a variety of traffic conditions.

replication — Copying the database from one server to another, as in the case of a WINS server.

replication interval — The amount of time between WINS pull replication requests.

replication trigger — The WINS pull partner's message that initiates replication with the push partner.

request — The DHCP client's acceptance of the DHCP offer.

resolver — A host that requests DNS name resolution.

reverse lookup zone (IN-ADDR.ARPA) — Useful for performing the reverse of a normal query: Instead of resolving a name to an IP address, it resolves an IP address to a name.

Router Information Protocol (RIP) — A distance vector-based protocol that identifies the best route for a destination based on the number of hops.

scope — A range of IP addresses that the DHCP server distributes to DHCP clients.

secondary domain servers — A DNS server that can receive a read-only copy of the zone database from a primary domain server. Secondary domain servers are useful for providing redundancy and load balancing.

Serial Line Internet Protocol (SLIP) — SLIP uses only the TCP/IP protocol and is useful for UNIX connections. SLIP is very basic and does not support authentication protocols, encryption, or compression.

Shiva Password Authentication Protocol (SPAP) — Shiva products (acquired by Intel) are a popular alternative to Microsoft RAS solutions. Shiva encrypts authentication credentials for Shiva LAN Rover software.

Simple Mail Transport Protocol (SMTP) — A mail protocol that transports only basic text. SMTP is the protocol that transfers or forwards mail to an email server.

Simple Network Management Protocol (SNMP) — A network monitoring and management protocol. Its usefulness is comprised of several elements that work together with the ultimate purpose of informing the administrator of a changing trend in the use of an object or alerting the administrator of an error, failure, or condition.

slave server — An authoritative DNS server that receives the zone transfer from the master and is named in the zone by an NS record.

SNMP agent — A service that runs on the actual object you want to monitor using an SNMP management system.

SNMP management system — Sends requests for information from the monitored system, known as the SNMP agent.

SONET — A fiber optic transmission medium that is self-healing. If a line is cut, traffic redirects to another ring.

Start of Authority (SOA) server — The authoritative server for information about the domain; the domain cannot function without it.

static IP address — An IP address that is manually and permanently assigned to a network host.

Systems Network Architecture (SNA) Server — A Microsoft product that acts as a gateway between the client and the mainframe.

time to live (TTL) — In DNS and WINS, the length of time a record is stored.

trap — A message issued from the SNMP agent to the SNMP management system.

turnkey fax servers — Self-contained, freestanding devices in which the software and hardware are already installed. Except for some company-specific configurations, they are ready to fax right out of the box.

virtual private network (VPN) — A communications session protected inside an encrypted virtual "tunnel" that is extremely difficult for intruders to breach. VPN is most commonly used over an Internet connection.

Windows Internet Naming Service (WINS) — A Microsoft NetBIOS name resolution service.

WINS proxy agent — A WINS-enabled computer that listens on the subnet for WINS broadcast messages, such as: query, refresh, release, and registration. The WINS proxy then communicates with the WINS server to resolve or register NetBIOS names.

zone transfer — A copy of the zone DNS database that is copied to another DNS server.

Review Questions

1. Which service automatically distributes IP configuration to clients?
   a. DHCP
   b. DNS
   c. SNA
   d. DISTIP
2. Which of the following is the correct order in the lease process?
   a. Offer, discovery, acknowledgment, request
   b. Discovery, acknowledgment, offer, request
   c. Discovery, acknowledgment, request, offer
   d. Discovery, offer, request, acknowledgment
3. DNS is an alternative to which text file?
   a. LMHOSTS
   b. HOSTS
   c. Readme.txt
   d. HOST
4. A resolver is:
   a. an authoritative server that solves conflicting IP address issues
   b. a service that removes outdated WINS records
   c. a server that returns iterative queries
   d. a client that requests DNS name resolution
5. A DNS zone is:
   a. a physical boundary where a host must reside in order to become a resolver
   b. a naming boundary
   c. another term for a Windows NT/2000 security domain
   d. the authoritative server for all name resolution within a domain

6. Which of the following DNS records resolves the host name to an IP address?
   a. PTR records
   b. SOA records
   c. MX records
   d. A records
7. Which of the following DNS servers does not have any permanent DNS records?
   a. caching-only
   b. slave
   c. secondary
   d. primary
8. Which operating system primarily uses NetBIOS names?
   a. Linux
   b. OS/2
   c. Windows
   d. NetWare
9. What does DDNS do?
   a. automatically registers NetBIOS names
   b. automatically registers DNS names
   c. recognizes when DNS server records are out of date and automatically replicates a more up-to-date zone
   d. dynamically configures a host's IP configuration
10. What does a push partner do?
    a. responds to pull partners' replication triggers and sends WINS updates
    b. responds to other push partners' replication triggers and sends WINS updates
    c. automatically replicates WINS data to DNS servers
    d. forces an update on pull partners
11. You have an OS/2 server that does not use WINS. The network is Windows NT 4.0. How can you make the OS/2 server available in WINS?
    a. use the WINS update agent
    b. use the WINS proxy agent
    c. use the SNMP agent
    d. You cannot do anything to resolve this.

12. Which of the following mail protocols allows you to search for specific mail on the server and download only select items?
    a. POP3
    b. SMTP
    c. MIME
    d. IMAP
13. What special hardware factors might you consider for an email server?
    a. additional hard disk space
    b. powerful processing capability
    c. additional RAM
    d. all of the above
14. Why might you want to have a dedicated, separate Internet connection to the web server? (Choose two.)
    a. So LAN traffic does not affect web server responsiveness.
    b. So Web traffic does not affect LAN responsiveness.
    c. This is a requirement for an effective firewall.
    d. A separate connection offers no practical benefit.
15. How might RAM improve web server performance?
    a. Web server software is particularly memory-hungry.
    b. RAM is not particularly an issue with web content.
    c. More web content can be cached into RAM and served faster than from the hard disk.
    d. FTP and streaming media can be stored wholly in memory.
16. What is one of the disadvantages of most UNIX services?
    a. The graphic-based interface is not consistent from one version of UNIX to another.
    b. The text-based administration is not intuitive.
    c. UNIX services are historically less stable than other operating systems' services.
    d. The UNIX services are completely incompatible with any other operating system.
17. What is the difference between PPP and SLIP?
    a. PPP offers a GUI, and SLIP does not.
    b. PPP supports more protocols than SLIP.
    c. PPP is available only on Windows systems, and SLIP is only available on UNIX clients.
    d. There is no perceptible difference besides the spelling.

18. How are fax servers better than fax machines? (Choose all that apply.)
    a. Fax boards transmit faxes much faster.
    b. Fax servers can better protect sensitive documents from prying eyes.
    c. Fax servers automatically fill in some information.
    d. Fax servers do not jam.
19. What is the purpose of an SNA server?
    a. connect Macintosh clients to the same network as Windows clients
    b. connect clients to mini and mainframe hosts over a terminal session
    c. separate the public Internet from the private LAN
    d. allow dial-up clients to access the LAN as if connected locally
20. You receive a message on your pager that a particular router is flooded. Which service made this notification possible?
    a. PAGER
    b. SNMP
    c. RAS
    d. PPTP

HANDS-ON PROJECTS

Project 9-1

In this project, you will configure a NetWare 5.1 server with the DNS and DHCP Management Console. This requires an existing NetWare 5.1 server and a Windows 2000 Professional or Server or greater client. If the Windows computer does not have Novell Client for Windows 2000, download and install it from *www.novell.com*. The NetWare 5.1 server should have a static IP address of 10.10.10.5 and a subnet mask of 255.255.255.0. The Windows 2000 client should have a static IP address of 10.10.10.1 and a subnet mask of 255.255.255.0.

1. Log on to the NetWare 5.1 server with Admin rights. Your instructor will tell you which account and password to use.
2. From the Windows 2000 client, ping the NetWare 5.1 server to verify connectivity. Enter **PING 10.10.10.5** at a command prompt. You should see replies come back.
3. Install the DNS-DHCP Management Console. Double-click **My Network Places**, and browse through **Entire Network**, **NetWare Services**, and **NetWare Servers**.
4. The name of the NetWare server should appear in a window. Double-click the NetWare server. If more than one appears, ask your instructor which one to double-click.
5. Double-click through the path **SYS:PUBLIC\DNSDHCP**.

6. Double-click the **Setup.exe** file to begin setup, and proceed through the setup prompts. At the end, you will be prompted to copy the snap-in files. Select this option and click **Next**.
7. You are prompted for the location of the NetWare Administrator utility. This is the location of NWAdmn32.exe. Use the Browse button to access the SYS:Public\win32 directory. Click **Next** to complete the steps in the wizard.
8. Browse to SYS:Public\Win32 and double-click the **NWAdmn32.exe** file. If a Tip appears, close it. (Consider making a shortcut to the desktop to make this more accessible.)
9. From the NetWare Administrator, click the **Tools** menu, then **DNS-DHCP Management Console**.
10. Leave the console open for the next project. Feel free to look around in the console, but do not change anything.

Project 9-2

In this project, you will access the DNS-DHCP Management Console and enter a DHCP scope. Then you will boot the Windows 2000 client to receive IP configuration from the DHCP server.

1. The DNS-DHCP Management Console is still open from Project 9-1. If it is not, refer to Steps 8 and 9 to open it.
2. At the top of the interface are two tabs: DNS Service (default) and DHCP Service. Click **DHCP Service**.
3. A single Our Network item appears on the left; there are no scopes created yet. On the toolbar, click the second icon from the left (it looks like a box).
4. The Create New DHCP Record appears. Select the **DHCP Server** item and click **OK**.
5. The Create DHCP Server text box appears. Type in the name of the NetWare 5.1 server in its context. Your instructor will tell you how to enter the context. You can also browse for it if you like. Then, click **Create**.
6. The server object appears at the bottom of the console. Click the "box" toolbar button again, and in the Create New DHCP Record dialog box, click **Subnet** and then **OK**.
7. The following information should be entered for servers that are not on the same network. If your class has multiple servers on the same network for this project, your instructor will provide different configuration information. For the Subnet Name, enter **ClassIPRange**, use the current NDS context, or enter a context specified by your instructor.
8. Enter a subnet address of **10.10.10.0** and a subnet mask of **255.255.255.0**. The dialog box should look similar to the one shown in Figure 9-19. When finished, click **Create**.

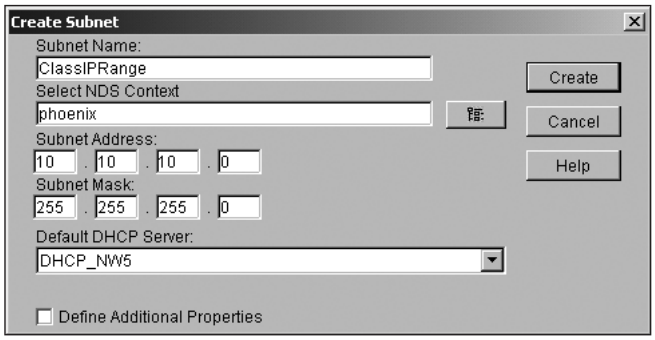

Figure 9-19 Creating a new DHCP record

9. The current DHCP configuration could cause a problem in a production environment. The NetWare 5.1 server IP address is 10.10.10.5, and yet this address is in the range you just created. You don't want a client to attempt to lease this address. (Even though there are mechanisms that prevent this, you don't want to take chances.) In the left pane, click the **10.10.10.0** (ClassIPRange) item and select the "Box" icon again. The Create New DHCP Record dialog box appears.
10. Click the **IP Address** item and click **OK**. Enter IP address **10.10.10.5**, verify that the Assignment Type is Exclusion, and click **Create**. Now, the DHCP server will not lease that IP address.
11. With the 10.10.10.0 (ClassIPRange) item still selected, click the **Other DHCP Options** tab. Then click the **Modify** button.
12. Configure the DHCP client's DNS server. Select the Domain Name Server (Code 6) and click Add.
13. At the bottom of the dialog box, click the **Add** button. Enter the IP address for an available DNS server as told by your instructor. Click **OK** to close all dialog boxes.
14. Boot the Windows 2000 computer and log on as a Domain Admin. Configure the IP address of the local Windows 2000 computer to use the newly configured NetWare DHCP server. Click **Start**, point to Settings, and click **Network and Dial-up Connections**. The Network and Dial-up Connections screen opens.
15. Your local area connection appears in the window. Right-click it and choose **Properties**.
16. Scroll down to Internet Protocol (TCP/IP), click it, and click the **Properties** button.
17. Click the **Obtain an IP address automatically** item and the **Obtain DNS server address automatically** item. Click **OK** to close all dialog boxes and reboot.

Windows 2000 should automatically request and receive a new DHCP-assigned IP address once you close the Local Area Connection dialog boxes. However, rebooting is a good way to ensure that a lease occurs.

18. After logging back on as a Domain Admin, click **Start**, click **Run**, type **CMD**, and press **Enter**.
19. In the Command Prompt window, type **IPCONFIG /ALL** and press **Enter**. You should see a 10.10.10.x address and a DNS server address as specified in Step 17. The DHCP server successfully leased an IP address to the Windows 2000 DHCP client.

Project 9-3

In this project, you will configure a Windows 2000 web server.

1. Log on to a Windows 2000 server using an account with Domain Admin membership. Your instructor will tell you what user name and password to use.
2. Using Windows Explorer, create a new folder named **My Web** on the C: drive.
3. Using Microsoft Word or any other word processor capable of saving a file as a web page, create a new document and type in a few words (careful, other people might see this later).
4. From the File menu, click **Save As** and navigate to the C:\My Web folder. Then from the Save as type drop-down list at the bottom, choose **Web Page** and name the document **Default**.
5. Click **Start**, point to Programs, point to Administrative Tools, and click **Internet Services Manager**.

Despite the fact that you are in a console called Internet Services Manager, you are using Internet Information Server 5.0.

6. In the left pane, you see the name of your server. Expand it and locate Default Web Site. Right-click it and click **Browse** to verify that IIS and the browser are operating correctly. Two browser windows open to an excellent online help resource about configuring a Windows 2000 web. Close the browsers.
7. In the left pane, right-click the server, then point to **New** and click **Web Site**.
8. The Web Site Creation wizard begins. Click **Next** to start configuration and after each step where it applies.
9. Enter a description of the site (Class Web, for example). This is for your reference only; the Internet public does not see it.
10. In the IP Address and Port Settings page, select the IP address of the local computer from the drop-down list. Click **Next**.
11. In the Path text box, enter **C:\My Web**. Be sure to leave the Allow anonymous access to this Web site checked.
12. Leave the default Web Site Access Permissions, and finish the wizard.
13. Right-click the new web site in the left pane and click **Browse**. Your web site now appears in the browser.

14. If connected to other web sites on the same network, you can ask another student for their IP address (use the IPCONFIG command to determine this) and type **http://IPAddress** where IPAddress is the other student's IP address. The other student's web site appears in your browser.
15. Close all open browsers and the Internet Information Services console.

Project 9-4

In this project, you will visit *www.inflow.com* to see an excellent example of not only a co-location facility, but also an enterprise server environment.

1. Using your web browser, access the *www.inflow.com* web site.
2. On the home page, there should be a link to Take a video tour. Click this item.
3. A RealPlayer streaming media presentation begins. If you do not have RealPlayer installed, follow the link to install RealPlayer and return to this site.
4. What kind of physical security does Inflow have?
5. When can you expect Inflow personnel to be available to monitor your applications and equipment?
6. What facility controls does the staff monitor?
7. How does Inflow ensure uninterrupted power?

Project 9-5

In this project, you will create a DNS zone in Windows 2000.

1. Click **Start**, point to Programs, point to Administrative Tools, and then click **DNS**. If this item does not appear, install it using the Add Remove Programs item in Control Panel. DNS is a subset of Networking Services.
2. Right-click the name of the server in the left pane, and choose **New Zone**. A wizard begins. Click **Next** to start it and after each step.
3. Select a standard primary zone.
4. Select a forward lookup zone.
5. Type a name for the zone. This would be the name of the zone only, not the name of existing or planned hosts. If you had a server named Mail1 and a domain named accusource.net, you would enter only accusource.net here, not Mail1.accusource.net.
6. A file stores the DNS database. Leave the default value and finish the wizard.

Project 9-6

In this project, you will create an A record in the zone created in Project 9-5.

1. In the left pane, select the zone you created in Project 9-5.
2. Right-click the zone and choose **New Host** from the menu. (You might have to refresh the zone by right-clicking on it, and then selecting "Refresh" from the menu.)

3. Enter any name here; this is a fictitious record.
4. Enter the IP address; any address will do.
5. Notice that Windows 2000 will also automatically create an associated PTR record if you so configure it. We haven't created a reverse lookup zone, so it doesn't apply here.
6. Click **Add Host**, and a confirmation screen appears. Click **OK** and then click **Done** on the New Host dialog box.
7. Open a command prompt, and ping the new record by its full name in the form: *host.domain*. For example, if you created student.class.com, type **ping student.class.com**. The request will time-out, but you will see that it is indeed attempting to ping the name and IP address you entered.
8. Delete the zone by right-clicking it and selecting **Delete**.

Project 9-7

In this project, you will create a RAS VPN server and a dial-up connection on a Windows 2000 server. Then you will connect to the VPN server.

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and click **Routing and Remote Access**.
2. Click on the name of the local server. Notice the red arrow indicating that the service is neither configured nor running. Right-click it and select **Configure and Enable Routing and Remote Access**.
3. Another handy wizard begins. Click **Next** to start it and again after each step.
4. Click on the **Virtual private network (VPN) server** option.
5. TCP/IP should be listed as a protocol. Click **Next**.
6. Specify the **<No internet connection>** option. Normally, you would use an Internet connection, but for our purposes, this should suffice.
7. Leave the Automatically setting for clients to receive IP addresses from a DHCP server.
8. Click **No, I don't want to set up this server to use RADIUS now**, and finish the wizard. After a few moments, the RRAS service starts.
9. Click the **Ports** item in the left pane. Notice that you have both PPTP and L2TP ports through which clients can connect. Because we have not configured a certificate structure, we will use PPTP.
10. Minimize the RRAS console.
11. Click **Start**, point to **Settings**, and click **Network and Dial-up Connections**. Double-click the **Make New Connection** item. Guess what, another wizard! Click **Next** here and after each step. (If you have not previously specified the area code in which the computer resides, you might be prompted to enter this information first.)
12. Click **Connect to a private network through the Internet**.

13. If the Windows 2000 computer has an existing dial-up connection, click **Do not dial the initial connection**. If we had a dial-up Internet connection, this screen would configure it to dial automatically when we initiated a VPN connection.
14. Enter the IP address of the local server. Again, you can use IPCONFIG from a command prompt to find this information.
15. It doesn't matter if you select For all users or Only for myself.
16. Leave the Enable Internet Connection Sharing for this connection unchecked, and finish the wizard.
17. The Connect Virtual Private Connection dialog box appears immediately. Type in your user name and password, and click **Connect**.
18. If the connection is denied, you will need to access your user account in Active Directory Users and Computers located in Administrative Tools (Start, Programs, Administrative Tools). In your user account properties, click the Dial-In tab and grant yourself access (see Figure 9-20 for reference). The default is to deny access as a security precaution.

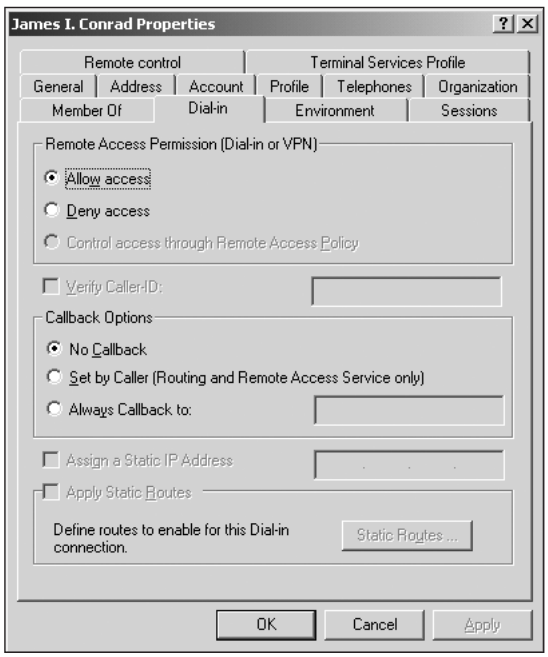

Figure 9-20 Grant yourself remote access permission if necessary

19. Besides practice for configuring RRAS and dial up connections, using a VPN to connect to the same local computer is an excellent way to troubleshoot other dial-in conditions and rules (compared to waiting a minute for a real modem to handshake and connect). Windows 2000 has complex but effective conditions under which users can dial up.

CASE PROJECTS

1. You have been hired to design a technology update for a medium-sized organization. The organization currently uses UNIX servers, Windows 2000 Professional clients, and a single Windows NT 4.0 server running SNA Server so that clients can run a terminal session to an AS/400 DB2 database. There are only about 200 users, but it is quite cumbersome for the administrator to manually enter the host records for each client and server into the DNS database file. Also, each desktop has statically configured IP information. The organization has made a significant investment in the AS/400 server and database, and you are required to keep it in the new plan, but you are authorized to purchase two or three new PC servers as needed. What should you propose to improve the administration of this network?
2. A local charity, KidHelp, asks you to contribute your knowledge and expertise. KidHelp is new to the Internet. KidHelp has been using a shared Internet connection through a cable modem for desktop client Internet access. Management is concerned that a malicious Internet user might attempt to somehow damage the KidHelp network servers or clients because the cable modem is always on. What can you recommend to KidHelp?